



Sponsor: OSD
Dept. No.: J82A
Contract No.: W15P7T-13-C-A802
Project No.: 0715D190EP

The views, opinions, and/or findings contained in this report are those of the authors and should not be construed as an official Government position, policy, or decision, unless designated by other documentation.

This technical data was produced for the U.S. Government under Basic Control No. W15P7T-13-C-A802, and is subject to the Rights in Technical Data-Noncommercial Items clause at DFARS 252.227-7013 (FEB 2012).

**Approved for Public Release;
Distribution Unlimited.
Case Number 15-3845**

©2016 The MITRE Corporation.
All rights reserved.

McLean, VA

Data Center Energy Efficiency Technologies and Methodologies

A Review of Commercial Technologies and Recommendations for Application to Department of Defense Systems

Performed at the request of the Assistant Secretary of Defense for Research and Engineering's Reliance 21 Energy & Power Community of Interest

**Alexander D. Schlichting, Ph.D.
November 2015**

Approved By



11/17/15

Kurt Eisenbeiser, Ph.D.

Date

Principal Power & Energy Systems Engineer

Task Leader for the Data Center Energy Efficiency Technologies and Methodologies Review

This page intentionally left blank.

Executive Summary

Although data centers (DC) are most often associated with warehouses filled with rack upon rack of servers owned and operated by large information technology (IT) companies, the Department of Defense (DoD) has a significant number of DCs themselves. While the DoD initiated a consolidation effort in 2010, there is still plenty of opportunity to improve the energy efficiency of their DCs to increase their cost effectiveness and resiliency. The goal of this study is to examine the technologies and related technology management methodologies commercial companies, hereafter referred to collectively as Industry, use to improve the energy efficiency of their DCs and identify the most promising ones that DoD DC operators could leverage. The study concludes with a set of key overall findings and a number of DoD-specific recommendations.

Key to understanding the DoD-specific recommendations are the unique DoD applications and requirements for their DCs. For the purposes of this study, the definition of a DC is extended beyond warehouses filled with hundreds of racks to any computing capability for local or remote use beyond those provided by a stand-alone desktop or hand held computing device. This introduces into the discussion a large number of mobile, tactical command, control, communications, and computer (C4) systems across the Services. A couple of examples are mobile command posts mounted on the back of an M1152 HMMWV, airborne surveillance and targeting platforms such as the E-8C JSTARS, and the combat direction center on an aircraft carrier. All of these platforms present unique challenges related to the integration of significant IT capabilities, and their operational effectiveness would benefit greatly from the IT system, or DC, being energy efficient. The other major DoD-unique requirements, used to inform this study's recommendations, relate to information security and access control. The overall Federal Government uses the FedRAMP certification program to ensure a minimum level of information security for Government DCs, and the DoD has leveraged this certification with its enhanced FedRAMP+ program. There are also additional information security requirements for different types of information as summarized by the Defense Information Systems Agency (DISA) DoD Cloud Computing Security Requirements Guide. Highly sensitive and collateral information requires additional access control and physical security on top of the FedRAMP+ requirements. DoD DCs can still make use of some of the Industry's technologies despite this combination of unique DC platforms and increased security requirements. DOD DCs also often present opportunities not present in the commercial realm.

A typical DC has the same major components no matter its application or size. The actual IT systems go beyond the IT computing systems, or servers. Data storage drives, networking equipment, and data switches are all considered part of the IT systems. The electrical power and cooling equipment is often collectively referred to as the DC physical infrastructure (DCPI). This term encompasses on-site backup generators, switchgear, uninterruptible power supplies (UPS), power distribution units (PDU), power cables, computer room air conditioners and handlers (CRAC/CRAH), chillers, direct expansion (DX) units, and even cooling towers. There are also a slew of additional miscellaneous components to a DC that need to be considered during its design: lighting, raised floors, exhaust plenums, and even the DC operators themselves. The Power Usage Effectiveness (PUE) is the typical Industry metric for determining the energy efficiency of a DC. It simplifies down to a ratio of the total amount of electrical energy consumed by the DC over the amount of energy consumed by the IT equipment alone. A value of 2 means that for every Watt dedicated to performing computations an additional Watt is required for the operation of the DCPI. An "ideal" DC would have a PUE of 1, where all of the

electrical energy is used for computations. This metric, however, focuses on the efficiency of the DCPI and assumes that the power consumed by the IT equipment is a constant amount that cannot be lowered. This is misleading because IT equipment has become more efficient over the past decade. The EPA has developed ENERGY STAR[®] certifications for many different types of IT equipment used in DCs, and purchasing certified equipment when available is one way of improving DC energy efficiency not captured in the PUE metric.

The technologies found during this review separate into three main categories: management methodologies, low-capital technologies, and high-capital technologies. In this report, an emphasis is placed on the management methodologies and low-capital technologies as they represent the paths of least resistance toward implementing more energy efficient DoD DCs. Table 1 summarizes all of the particularly promising technologies and management methodologies within each of the three categories mentioned above. Table 2 summarizes the ones that would be a good fit for a specific, DoD unique, application. Table 3 gives an overview of some of the other technologies investigated that were not considered technologically mature enough for immediate adoption. After the tables are summaries of the key findings contained within this report and the DoD-specific recommendations that resulted from the study.

Table 1. A summary of the promising DC energy efficiency technologies for DoD unique applications.

Technology	Pros	Cons
<u>Containerized, pre-fabricated DCs</u>	<ul style="list-style-type: none"> • Very mobile with fast deployment as they are typically in ISO shipping containers but are also available as individual racks • Pre-fabricated and configured to operate efficiently by the vendor and can be purchased with the most recent technology • Allows for easy right-sizing of IT capacity for efficient operation 	<ul style="list-style-type: none"> • Difficult to repair on-site • Requires ability to provide necessary power and water hook-ups that can become expensive at large scales • Can be wasteful if a single unit is more IT power than a mission requires • Expensive on a large scale compared to purchasing just the individual servers, racks, etc. as the vendor charges for each container to be assembled and configured • Having both containerized IT and DCPI on a large scale can promote purchasing excessive DCPI to provide redundancy to each IT container individually, as opposed to using networked DCPI in N+1 configurations
<u>Free cooling using surrounding environment</u>	<ul style="list-style-type: none"> • Use of Free Cooling from readily available sources, such as seawater or the air surrounding a plane, through an economizer mode significantly decreases annual cooling energy requirements 	<ul style="list-style-type: none"> • Requires necessary infrastructure to transport seawater or air to the DC, increasing the energy consumed by fans and pumps
<u>Solid State Drives (SSD) for mobile and tactical applications</u>	<ul style="list-style-type: none"> • Significantly lower cooling load • Higher maximum safe operating temperatures • Significantly faster ramp-up and ramp-down times • Higher tolerance for vibrations 	<ul style="list-style-type: none"> • Cost more than Hard Disk drives, however combination solutions are available with both SSDs and Hard Disks to reduce the cost compared to an all SSD solution
<u>Co-Generation and Tri-Generation</u>	<ul style="list-style-type: none"> • Uses waste heat from prime power generation to power absorption chillers and reduce electrical energy demand of DC • Multiple DoD platforms have significant on-site prime power generation systems 	<ul style="list-style-type: none"> • Requires significant capital expense to route prime power waste heat to onboard DC • Other applications, such as onboard desalination, might already be using the prime power plant waste heat • Requires a substantial minimum load to maintain energy efficiency
<u>Waste heat for desalination</u>	<ul style="list-style-type: none"> • DC waste heat can provide necessary heating power for certain desalination systems, significantly increasing energy efficiency of combined IT and desalination systems • Can operate with waste heat qualities (temperatures) available from liquid and two-phase cooling systems • A multiple stage hybrid absorption heat pump desalination system can 	<ul style="list-style-type: none"> • Not a mature technology with few commercial companies advertising in this area but it could be leveraged for unique DoD applications

Technology	Pros	Cons
	completely remove the DC heat and replace a traditional chiller	
Direct-to-chip cooling	<ul style="list-style-type: none"> • Highest cooling performance levels by minimizing distance between heat source and cooling fluid, thereby maximizing heat transfer • Methods exist that combine both high performance cooling fluids and high heat flux phase change cooling 	<ul style="list-style-type: none"> • Requires specialized IT equipment and infrastructure to deliver cooling fluid to each server • High capital cost • Less mature technology mostly focused on research platforms but it could be leveraged for unique DoD applications • Compared to air cooling, liquid cooling has a greater chance of failure due to leaks

Table 2. A summary of the pros and cons of all of the promising technologies and management methodologies researched, focusing on large-scale DC facilities.

Technology	Pros	Cons
Management Methodologies		
Management and Monitoring	<ul style="list-style-type: none"> Some DCs already have a number of distributed sensors installed due to the real-time, integrated data reporting requirements for ENERGY STAR® certified servers Informed predictions of the impact of specific DC energy efficiency improvements can be made using distributed sensor data and computational fluid dynamics (CFD) simulations 	<ul style="list-style-type: none"> There is no standard format for servers reporting their power consumption and temperature data as ENERGY STAR® certification does not specify one, collecting the data can become challenging in a diverse IT environment Fully-integrated DCIM is expensive and has a history of not delivering on its promises or hype DC CFD simulations are limited in their accuracy by not only the quality of the information used to run the simulation but also by their applicability to only a single instant in time
Right-sizing throughout data center lifetime	<ul style="list-style-type: none"> Multiple methods available: containerized systems, section-by-section, or rack-by-rack Distributes capital costs over a DC's lifetime and decreases overall capital costs Increases overall DCPI energy efficiency by having installed capacity more closely match current IT load 	<ul style="list-style-type: none"> Containerized DCs or DCPI increase difficulty of on-site maintenance Containerized DCs and DCPI require facility infrastructure to provide necessary power, water, and other hookups as new containers are installed Section-by-section modularization of a DC would likely require a concurrent upgrade to close-coupled cooling Energy efficiency gains disappear for DCs with a known, fixed IT load throughout their lifetime, particularly if the lifetime is only a few years
Decommissioning, upgrading, and consolidating	<ul style="list-style-type: none"> Newer servers are more energy efficient at a given computational load Flash data storage produces a lower cooling load, has shorter ramp-up and ramp-down times, and can operate safely at higher temperatures than hard disk systems Blade servers increase the opportunities for waste heat reuse with higher heat qualities 	<ul style="list-style-type: none"> Applications on legacy systems need to be modernized Blade servers increase the cooling load density and often require an upgraded cooling system for widespread implementation Upgrading applications to modern operating systems and replacing IT equipment has a high capital cost Flash data storage is expensive, but hybrid drives are available to mitigate capital costs
Virtualization for flexibility	<ul style="list-style-type: none"> Disassociation of applications from specific physical servers requires fewer overall servers to support a given IT load at the desired resiliency/availability Increases average server utilization level, increasing their level of computational performance per unit 	<ul style="list-style-type: none"> Higher individual server utilization levels increase heat gain density Consolidating remaining servers to fewer racks even further increases heat gain density Major energy efficiency benefits of virtualization cannot often be

Technology	Pros	Cons
	<p>electric power, or energy efficiency</p> <ul style="list-style-type: none"> Some virtualization software tools have the capability to power on and off specific servers based on current IT needs, eliminating unnecessary idle server power consumption 	<p>realized without also upgrading cooling system to accommodate higher density heat loads</p>
Low-Capital Technologies		
Air Flow Management and Hot/cold containment for air cooling	<ul style="list-style-type: none"> Range of options available for both Air Flow Management and Hot/Cold Containment, both can be low cost, flexible, and quick to install while still effective Significantly increases cooling efficiency and produces greater inlet temperature consistency Increases maximum per-rack IT load cooling system can handle 	<ul style="list-style-type: none"> Standard methods use raised floors and overhead exhaust plenums, however not required Most advanced systems can be expensive to install, however there are effective low-cost options Need to decrease cooling system output by increasing server inlet temperature or cooling unit set points to realize efficiency gains
Economizer modes for chillers	<ul style="list-style-type: none"> Use of Free Cooling through an economizer mode significantly decreases annual cooling costs Decoupled economizer modes operating simultaneously with traditional cooling system increase cooling system efficiency Some methods make use of high performance evaporative cooling 	<ul style="list-style-type: none"> Use of outside air requires larger fans to move the air Direct use of outside air requires additional filtering and humidity control, however there are effective indirect methods Some methods require additional cooling fluid loops and the use of a cooling tower
Power path optimization	<ul style="list-style-type: none"> Upgrading to modern uninterruptible power supplies (UPS) and power distribution units (PDU) can significantly increase power system efficiency Switching to a more Vdc-centric form of facility power distribution can decrease required DC floorspace by reducing the number of discrete systems 	<ul style="list-style-type: none"> For DCs using utility grid electricity the overall conversion step will remain the same (i.e., 480 Vac to 1 Vdc) down to the processors Conflicting reports cast doubt on whether upgrading to a Vdc-centric system can result in reduced operating expenses from increased power system efficiency compared to a modernized traditional system
Supplemental evaporative cooling/adiabatic cooling (e.g., foggers)	<ul style="list-style-type: none"> Can be added to traditional air-cooled DCs or economizer modes that make use of outside air Inexpensive and simple way of decreasing cooling and potentially humidity control costs 	<ul style="list-style-type: none"> Potential to decrease DC cooling costs very dependent on external environment and therefore DC location Can make it more difficult to keep IT environment within its recommended humidity range
Close-coupled cooling	<ul style="list-style-type: none"> Cooling efficiency significantly increased by decreasing the distance between the heat source and cooling system Modularity allows cooling capacity to be added as it is needed and more precisely controlled to improve efficiency through better load matching Rear door heat exchangers can bring high performance cooling fluids directly to the rack 	<ul style="list-style-type: none"> Smaller fans or cooling systems have lower maximum efficiencies than larger systems when both are at full utilization Initial capital cost for upgrading can be relatively high Facility needs to be flexible enough and be outfitted to allow for distributed cooling systems Humidity control is important, especially for rear door heat exchangers, as excess moisture can

Technology	Pros	Cons
	<ul style="list-style-type: none"> Modularity combined with increased cooling efficiency and power can accommodate higher density, localized, heat gains 	condense in close proximity to the IT equipment
High-Capital Technologies		
Direct-to-chip cooling	<ul style="list-style-type: none"> Highest cooling performance levels by placing cooling liquids into direct contact with the heat sources Methods exist that combine both high performance cooling fluids and high heat flux evaporative cooling 	<ul style="list-style-type: none"> Requires specialized IT equipment and infrastructure, especially rack piping, to deliver cooling liquids to each server High capital cost and not widely commercially available Less mature technology mostly focused on research platforms
Waste heat reuse	<ul style="list-style-type: none"> Can increase overall site energy efficiency by combining DC facility with local space heating needs Potential way to share the capital expense of upgrading a DC through agreements with heat providers and consumers Has the potential to drive water desalination systems 	<ul style="list-style-type: none"> Dependent on local need for low quality heat for most DCs High quality heat requires more expensive liquid and two-phase (evaporative) cooling systems, ideally direct-to-chip systems Limited proven applications beyond space heating or providing hot water
Co-Generation and Tri-Generation	<ul style="list-style-type: none"> Uses waste heat from prime power generation to power absorption chillers and reduce electrical energy demand of DC Increases DC resiliency by eliminating its reliance on vulnerable utility grids More efficient than grid-tied systems due to both decreased power transmission losses and use of waste heat 	<ul style="list-style-type: none"> Requires significant capital expense to install on-site power generation, however positive ROIs are expected

Table 3. Nascent technologies potentially worth R&D investment or following

System Efficiencies	Description
<u>DC waste heat to power desalination systems</u>	Uses the waste heat produced by the DC IT equipment to power a desalination plant. The desalination process could also have the additional benefit of completely cooling the water such that it doesn't require a mechanical chiller or other cooling before it is sent back to the DC.
<u>Instrumenting DCs and using autonomous management software for efficient operations</u>	Distributed temperature, airflow, and power sensors can provide comprehensive and detailed information on DC energy efficiency performance. With this information, control programs can autonomously reroute IT tasks and adjust DCPI system parameters concurrently to maximize the real-time energy efficiency for the current IT load.
<u>DC modeling combining thermal simulations with system design tools</u>	Current DC models consist of either CFD room airflow simulations or high level system models using component specifications, however both are needed to accurately predict the performance of a DC design. There are ongoing efforts within ASHRAE TC 9.9 to create integrated, comprehensive, DC models.
<u>Direct-to-Chip cooling, including Two Phase cooling</u>	Most systems are currently in the research testbed phase due to the significant modifications required to the DC hardware. However, because these methods have high potential thermal fluxes they will likely continue to be developed.
<u>Thermoelastic cooling</u>	This process uses the latent heat of the solid-to-solid martensitic phase change in thermoelastic materials, such as Shape Memory Alloys, to pump heat. This is one of the multiple solid-state methods (also Magnetocaloric and Electrocaloric Cooling) proposed to replace vapor compression cycles. The primary benefit is the elimination of environmentally harmful refrigerant vapors.
<u>Piezoelectric infrasonic and ultrasonic fans</u>	Solid-state piezoelectric materials can actuate cantilevers to produce convection cooling that is more efficient than traditional axial systems at very small scales. The technology is currently limited by the resiliency of piezoelectric materials and the need for high voltage and frequency actuation signals.
<u>Jet impingement</u>	Directing high velocity streams of a fluid at a surface results in very high heat transfer coefficients between the fluid and the surface. The major challenges are the trade-off between the jet diameter and the heat transfer coefficient in addition to fluid management at the micro scale.
<u>3D stacked chip architectures</u>	This processor topology commonly puts a reconfigurable topology on top of a manycore topology and uses a software virtualization layer on top of the base kernel to allow for self-adaptation and dynamic reconfiguration based on the task at hand. The real-time flexibility allows for the optimization of load balancing, power consumption, and resilience down to the individual processor level.
<u>Heterogeneous System Architectures (HSA)</u>	This effort aims to optimize the power efficiency of a graphics processing unit (GPU) and central processing unit (CPU) pair by seamlessly sharing data and distributing tasks between the two.
<u>High temperature electronics</u>	Wide bandgap materials can operate at significantly higher temperatures and more efficiently than conventional electronics materials. The ability to run DC IT equipment at higher temperatures can significantly improve the efficiency of DC cooling by increasing the range of environmental conditions when free cooling is available and increasing the efficiency of traditional mechanical cooling.
<u>Reversible Logic</u>	Conventional computing results in wasted energy in the form of waste heat because logic gates produce multiple garbage states unrelated to the current calculation and erasing a bit of information always dissipates a small amount of heat. Reversible logic computing aims to both eliminate the garbage states and eliminate the need to erase bits of information with logically and physically reversible gates, which requires the potential outputs to be a one-to-one mapping with the inputs. The challenge has been developing a physical logic gate capable of operating in this manner and the vast majority of the work in this area is still theoretical.

Key Findings:

- Existing DoD DC software applications and IT hardware should be inventoried, consolidated, and modernized as the first step toward a more energy efficient DC. Legacy and/or idling servers have traditionally been the biggest factor contributing to DC energy waste. Increasing the average server utilization level (the portion of its maximum computing capacity actively in use) through consolidation, and potentially through virtualization, will significantly increase a DC's overall energy efficiency if the DCPI can accommodate the increased density and the now empty portions of the DC can either be used to increase overall computing capacity or isolated.
- In addition to improving PUE, the metric widely used by Industry to report a DC's energy efficiency, DCs should aim to maximize the energy efficiency of IT hardware, particularly servers. PUE does not incentivize decreasing the power consumption of IT hardware, so separate efforts are needed. ENERGY STAR[®] certification or the PAR⁴[®] energy efficiency evaluation tool can be used to motivate the sourcing of energy efficient servers.
- Airflow management and hot/cold air containment can be performed in an effective, modular, and low-cost manner, leading to significant energy efficiency gains for many DoD DCs. DC cooling systems are the second largest component of overall electric energy consumption after the IT hardware, so all types of DCs, from tactical to installation, should use some version of the available technologies.
- Right-sizing, or only adding DC IT and its associated DCPI when it is needed, can drastically increase the energy efficiency of a DC compared to one where all of the envisioned capacity is installed upfront. In addition to reducing operating expenses by lowering electrical power consumption, right-sizing also lowers the overall lifetime capital costs of a DC by reducing or even eliminating unnecessary equipment purchases.
- Free Cooling, the process of taking advantage of low temperature air or water in the environment, should be investigated for all DoD DC applications, not just those in cold environments. Decoupled economizer modes for DC cooling systems allow Free Cooling to supplement mechanical cooling systems and increase the overall cooling system efficiency.

DoD-Specific Recommendations:

- Encourage vendors to develop and produce low-cost, energy efficient servers by requiring the use and reporting of a related Industry standard, such as requiring ENERGY STAR[®] certification or the reporting of results from SPEC's SERT tool or from the PAR^{4®} tool managed by Underwriters' Laboratories and Tier44.
- Investigate low-cost airflow management and hot/cold containment solutions for all DCs, regardless of size or mission.
- Use section-by-section modularity for the large-scale facilities resulting from DoD DC consolidation efforts.
- Investigate containerized (either in ISO shipping containers or in individual all-inclusive racks) modular systems to determine if they can meet survivability specifications for DoD missions requiring fast deployment and/or high mobility.
- Ensure that DCs handling Impact Level 5 or 6 information and all tactical DCs are virtualized to increase the resiliency, reduce the footprint, and significantly reduce the cost of either fielding the systems or maintaining systems with the required levels of access control (DISA Cloud Computing Security Requirements Guide: 5 is the highest level of controlled, UNCLASSIFIED information and 6 is all CLASSIFIED information up to the SECRET level).
- Virtualized large-scale DC facilities should take advantage of existing software tools that automatically power on and off individual servers based on the current IT demands, minimizing unnecessary idle server power consumption.
- Replace traditional Hard Disk data storage for mobile and tactical applications with Solid State Drives (SSD). SSDs have significantly lower cooling loads, higher maximum safe operating temperatures, and significantly faster ramp-up and ramp-down times. Coupled with their ability to operate in higher temperature environments, SSDs have a higher tolerance for vibrations and therefore would perform better in many mobile and tactical environments.
- Implement CoGen systems with on-site power generation for all mission critical, large-scale DCs to significantly increase their resiliency by isolating them from the utility grid in addition to increasing their energy efficiency.
- Investigate using Free Cooling on large-scale naval vessels and high altitude, airborne platforms. Those platforms are not typically designed to consider the energy efficiency of the onboard IT or C4 equipment, however many operate in environments with significant opportunity for Free Cooling from the surrounding water (naval vessels) and air (airborne platforms).
- Utilize large platform prime power waste heat to power absorption chillers in CoGen or TriGen configurations for onboard DC cooling.
- Investigate whether DC waste heat can be used to supplement the platform power plant waste heat for desalination systems on naval vessels.

This page intentionally left blank.

Acknowledgements

This report was funded by, and performed at the request of the Assistant Secretary of Defense for Research and Engineering's Reliance 21 Energy & Power Community of Interest.

During the research for and writing of this report, the author received assistance from many different colleagues throughout MITRE and the U.S. Government and he would like to take this opportunity to thank them for their help and encouragement. Without naming them all, he extends his sincere gratitude to Mr. Henry Amistadi in particular, whose experience and insight into the field of data center design and operation contributed significantly to the quality of the information and the recommendations presented.

Table of Contents

Key Findings:	xii
DoD-Specific Recommendations:	xiii
1 Introduction	1
1.1 DoD Needs	1
1.2 Typical Data Center Components, Energy Consumption, and Power Losses	3
1.2.1 Data Center Components	3
1.2.2 Data Center Metrics	6
1.2.3 Data Center Design Best Practices for Energy Efficiency	9
1.3 Technology Taxonomy	10
2 Management Methodologies	12
2.1 Data Center Floor Monitoring and Management	12
2.2 Right-Sizing throughout Data Center Lifetime	14
2.3 Decommissioning and Upgrading	16
2.4 Virtualization for Flexibility	18
2.5 Summary and Recommendations	19
3 Low-Capital Technologies	21
3.1 Airflow Management	21
3.1.1 Virtual Containment	23
3.1.2 Hot/Cold Containment	24
3.2 Economizer Modes for Chillers	27
3.2.1 Dry Side	29
3.2.2 Wet Side	30
3.3 Power Path Optimization	31
3.4 Evaporation Cooling using Foggers	33
3.5 Close-Coupled Cooling	34
3.6 Summary and Recommendations	35
4 High-Capital Technologies	37
4.1 Direct-to-Chip Cooling	37
4.2 Waste Heat Reuse for Space Heating	38
4.3 Co-Generation and Tri-Generation	39
4.4 Summary and Recommendations	42
5 Nascent Technologies	42
5.1 Filtration and Desalination	42
5.2 Two-Phase Direct-to-Chip Cooling	43

5.3	Immersion Cooling	43
5.4	Autonomous Data Center Management.....	44
5.5	Integrated Data Center Thermal and System Modeling	44
5.6	Thermoelastic, Magnetocaloric, and Electrocaloric Cooling	45
5.7	Piezoelectric Fans	45
5.8	Jet Impingement.....	46
5.9	3D Stacked Chip Architectures.....	46
5.10	Heterogeneous System Architectures (HSA).....	46
5.11	High Temperature Electronics	47
5.12	Reversible Logic	47
6	Conclusions	48
7	References	50

List of Figures

Figure 1. Schematic of the typical power distribution network in a DC.	4
Figure 2. Schematic showing the flow of the heat out of the typical DC.	5
Figure 3. Schematic of the typical DC demands and their losses.	6
Figure 4. Distribution of DC power consumption.	7
Figure 5. Technology taxonomy of data center energy efficiency technologies.	11
Figure 6. Example display for monitoring DC state and performance.	13
Figure 7. Effect of right-sizing on designed and actual DCPI capacity over DC lifetime.	14
Figure 8. Distribution of the ultimate DC utilization levels for a typical DC.	14
Figure 9. PUE of a typical DC as a function of the utilization level.	15
Figure 10. TCO breakdowns for a modular DC using prefabricated containers.	16
Figure 11. Comparison between traditional and virtualized DCs.	18
Figure 12. Hot-aisle containment of DC with a raised floor and exhaust plenum	22
Figure 13. CCF for different evaluated DCs.	23
Figure 14. Proper airflow management without increasing cooling setpoints.	24
Figure 15. Rack-level exhaust air chimney system.	25
Figure 16. Vertical cooling configuration.	25
Figure 17. A magnetically attached modular containment system.	26
Figure 18. Cold-aisle containment without decreasing the cooling unit set points.	27
Figure 19. Number of hours per year a dry-side economizer mode could operate.	28
Figure 20. Number of hours per year a wet-side economizer mode could operate	28
Figure 21. The effect of an economizer mode on the PUE of a DC	29
Figure 22. Indirect air-side economizer with supplemental evaporative cooling	30
Figure 23. Schematic of a heat wheel used to provide free cooling to a DC.	30
Figure 24. Wet-side economizer mode based on evaporative cooling	31
Figure 25. A traditional 480 Vac to 208 Vac power distribution chain for a DC.	32
Figure 26. Row-based Vdc UPS used to provide -48 Vdc server input.	33
Figure 27. An evaporation cooling fogging system.	34
Figure 28. A direct-to-chip cooling prototype system with waste heat reuse.	37
Figure 29. Blade server with direct-to-chip cooling.	38
Figure 30. Schematic of a single-effect absorption cycle.	40
Figure 31. Schematic of a DC TriGen system.	41
Figure 32. Schematic of the FO desalination process.	43

List of Tables

Table 1. Promising DC energy efficiency technologies for DoD unique applications.....	vi
Table 2. Promising technologies and management methodologies for large DC facilities.	viii
Table 3. Nascent technologies potentially worth R&D investment or following.....	xi
Table 4. DISA Cloud Computing Security Requirements Guide information impact levels.	2
Table 5. Promising technologies or management methodologies.	11
Table 6. Waste heat quality available using main DC cooling techniques.	38

This page intentionally left blank.

1 Introduction

The use of data centers (DC) has been growing globally such that some market analysts are predicting the global DC power market alone could reach over \$23 billion by 2019 [1]. This estimate is partially spurred by the revelation that from 2011 to 2012, global DC energy consumption increased by 19%. In North America this rate slowed to a 6.8% increase in power consumption in 2013. Even so, DCs are still estimated to consume over 100 billion kWh of energy per year in the U.S., more than 2% of overall electricity use [2]. Part of this change in course can be attributed to the realization that better DC management is needed, as a 2014 survey by the Wired Real Estate Group showed 90% of US companies overprovision DC physical infrastructure (DCPI) by 25 – 50%, significantly increasing DC initial capital and ongoing operational costs [3]. However, there has also been a significant push to improve the energy efficiency of the DC hardware itself. This is exemplified by the addition of DCs to the Department of Energy’s (DOE) Better Buildings Initiative earlier this year, with the Defense Information Services Agency (DISA) being one of the notable participants [2], [4]. The purpose of this report is to examine the technologies and methods currently implemented by major U.S. companies, hereafter referred to as Industry, that can potentially be leveraged by the Department of Defense (DoD) to increase the energy efficiency of their DCs.

1.1 DoD Needs

In February 2010, the Federal CIO Vivek Kundra established the Federal Data Center Consolidation Initiative (FDCCI) to reduce the expenses associated with the growing Federal Government IT needs. The Secretary of Defense followed up a couple of months later, announcing the DoD data center consolidation effort. The DoD strategic goals for this effort are to reduce cost, reduce environmental impact, improve efficiency and service levels via automation, and enhance business agility and effectively manage change [5]. The intended impacts of these goals are to enhance mission effectiveness, improve security, and streamline IT provisioning and effectiveness. The Services and Components each have their own stated DC reduction goals, however a common theme is the virtualization and migration of enterprise-level services and applications to DISA DoD Enterprise Computing Centers (DECC). Preliminary estimates show a potential savings of \$1 billion annually by FY16, of which \$58 million comes from reduced energy costs. The DoD Chief Information Officer (CIO) Terry Halvorsen has also instructed the Services and Components to focus on the cost of performance as opposed to just the number of servers, eliminate redundant and outdated applications, and consolidate similar backend databases [6].

The DoD is restricted in its efforts in ways commercial companies and even other Federal Government Agencies are not, however. The DoD Cloud Computing Security Requirements Guide from DISA describes the requirements that set it apart [7], detailing the specific access and protection needs for DoD Information Technology (IT) systems in accordance with the Federal Risk and Authorization Management Program Plus (FedRAMP+) used by the DoD. FedRAMP+ certification for IT service companies, such as cloud service providers, is intended to allow for fast and cost-effective acquisition of IT services by the DoD without compromising the necessary security. Table 4 describes the different information impact levels and their associated DC requirements. Although the hosting of levels 2 & 4 information could potentially be outsourced to commercial companies, and such efforts are being explored [6], levels 5 and 6

plus information with higher classifications levels than SECRET must be effectively and efficiently handled by the DoD, itself.

Table 4. A summary of the DISA Cloud Computing Security Requirements Guide information impact levels [7].

Information Impact Level	Description	DC Requirements (on top of those already stated for lower levels)
1	Unclassified information approved for public release; merged with level 2	N/A
2	Non-controlled unclassified information: information approved for public release plus some DoD private unclassified information which requires minimum access control	If not physically located on DoD premises, must be within 50 states, D.C., and outlying areas of U.S.
3	Controlled unclassified information; merged with level 4	N/A
4	Controlled unclassified information: requires protection from unauthorized disclosure, mission critical data, export controlled data, private information, protected health information	Same as level 2
5	Controlled unclassified information: any information deemed requiring a higher level of access control than provided by level 4	Must reside in a facility under the exclusive legal jurisdiction of the U.S. Physical separation from non-DoD/non-Federal tenants is required.
6	Classified information up to SECRET	Stored and processed in dedicated infrastructure approved for classified information. Hosting organization must comply with NISPOM

To fully realize the diverse array of IT service needs and DC specifications, the definition of what constitutes a DC is expanded in this report beyond the colloquial norm: any computing capabilities for local or remote use beyond those provided by a stand-alone desktop or portable computing device. This ranges from large facilities with hundreds of racks filled with servers to single racks or blade enclosures. Large-scale facilities are the typical format of choice for Industry and the DoD has many analogous facilities; however, the DoD also has many different DCs at the other end of the spectrum. An example DC used by ground forces is the Northrop Grumman Command Post Platform Rigid Wall Shelter (CPP-RWS), which is mounted on the back of an up-armored M1152 HMMWV [8]. The CPP-RWS contains customizable racks for Command and Control (C4) equipment. It also has its own backup generator, power distribution, and environmental control. Some other examples of DoD-specific DC requirements can easily be imagined: the Air Force flies C4 equipment on many different airborne platforms and the Navy does the same with its vessels. The USAF E-8C JSTARS platform is a modified Boeing 707-300 with the necessary radar and computer subsystems to perform its battle management, command and control, intelligence, surveillance and reconnaissance missions [9]. Investigating new energy efficient ways to integrate and operate the radar and computer subsystems is imperative as a new competition is underway to develop its replacement [10]. Also, aircraft carriers must be able to support a similar mission for Navy fleets, controlling operational pictures of air, surface, undersea, strike, and even information warfare all from the carrier's combat direction center [11]. These types of DCs have few, if any, analogs in the commercial sector. When these platforms are

designed, the energy efficiency of the IT capabilities themselves is often an afterthought. This means that while the original platform may perform its principal mission in an effective and efficient manner, when a DC is integrated that is likely no longer the case.

There are significant opportunities for energy efficiency improvement in many DoD DCs, even the DoD's large-scale facilities. Since many Industry technologies can be directly applied to DoD's large scale facilities, the bulk of this report focuses on Industry's efforts to increase the energy efficiency of their large-scale DCs and the technologies the DoD could potentially leverage for their own. Some of these technologies could also be adapted to improve the efficiency of smaller-scale DoD DCs. This report also highlights the technologies that warrant further investigation they are likely candidates for adaptation to unique DoD needs. Lastly, it is noted that in Industry the common goal is to reduce the cost of a given capability level. The DoD's common goal is to increase mission effectiveness and overall capabilities within a given budget. Although two different goals at face value, DC energy efficiency helps both the DoD and Industry down their own respective paths, and there are significant opportunities for the DoD to leverage Industry IR&D efforts.

1.2 Typical Data Center Components, Energy Consumption, and Power Losses

Although there are many different specific types of DCs, this section gives a brief overview of the different components, their respective energy consumption, and the biggest sources of power loss (i.e., inefficiency) in a typical DC. Again, the focus is on the standard DC facility in use by both Industry and the DoD.

1.2.1 Data Center Components

The principal purpose of a DC is to provide IT computing power, which can be performed using a variety of form factors that all fall under the category of servers. However, the basic building block for larger DCs is the rack, the standard being the 42U configuration [12]. Its internal dimensions are 19 inches wide and 73.5 inches tall with 1.75 inches allocated per unit, or "U." The depth, external dimensions, and features vary by both model and manufacturer. The main two types of servers are rack servers, which can be 1U, 2U, or "xU" high, and blade servers [13]. Almost every watt of electrical power that goes to the servers is converted directly into heat, so the standard assumption is that the IT power and the resulting IT cooling load are equal. Modern rack servers typically have their own individual cooling fans and have significant expansion possibilities for adding additional storage drives. They also can support four or more multi-core processors and a significant amount of random-access memory (RAM).

Blade servers condense the computing power and bare necessities into a more compact package that is mounted, often vertically, in a separate chassis mounted in the standard 42U rack. The chassis takes over many of the connectivity and auxiliary functions that would otherwise come with each rack-style server (e.g., keyboard, video, mouse (KVM) connections, power distribution, and localized cooling fans). Sometimes blade chassis also contain networking switches, which would otherwise need to be a separate rack-mounted piece of equipment. Compared to rack servers, blades have lower localized storage capabilities and individual unit computing power. Also, proprietary blade-to-chassis connections restrict or prohibit the

compatibility one vendor's blade server with another's blade chassis, potentially causing equipment supplier issues.

Besides the servers, there are two other main types of IT equipment in a DC: networking equipment and data storage. Whether using localized hard drives (or flash drives) in rack servers or network-accessed data storage in its own rack-mountable configuration [14], data storage is a necessary and often overlooked component of a DC. It can represent a non-negligible amount of the DC power consumption and cooling load, particularly hard disks (which also often have lower operating temperature requirements than the rest of the IT equipment). In addition to the data storage, there is a diverse array of networking equipment to manage the data flow, protect the network, etc. Some examples are firewalls, intrusion prevention/detection systems, secure remote access systems, Ethernet switches, and more. These systems can come in 1U, 2U, all the way up to 20U configurations, as with the HP 10500 Series modular Ethernet switches [15].

Powering and cooling all of the IT equipment requires an array of equipment itself, collectively referred to as the DCPI. Figure 1 shows the typical power distribution network in a DC. First, even if the DC pulls its electricity from an electric utility grid, there are often on-site backup diesel generators to maintain operations in case of grid failure or long interruptions. The DC can switch from the utility grid to the backup generator in the case of grid failure using an automatic transfer switch (ATS). For shorter-term interruptions or when the grid electricity quality is poor, an uninterruptible power supply (UPS) is used to protect the DC equipment by providing a constant, high quality, source of power. The UPS either has an integrated energy storage backup, such as an array of large-scale batteries, or an external backup connected to it. Next, power distribution units (PDUs) allocate power to the different sections of the DC IT, for example, to each rack. The two pieces of switchgear, shown in Figure 1, often serve as power converters in addition to locations where the circuit can be disconnected manually. Their presence and exact function varies based on the specific equipment in a DC.

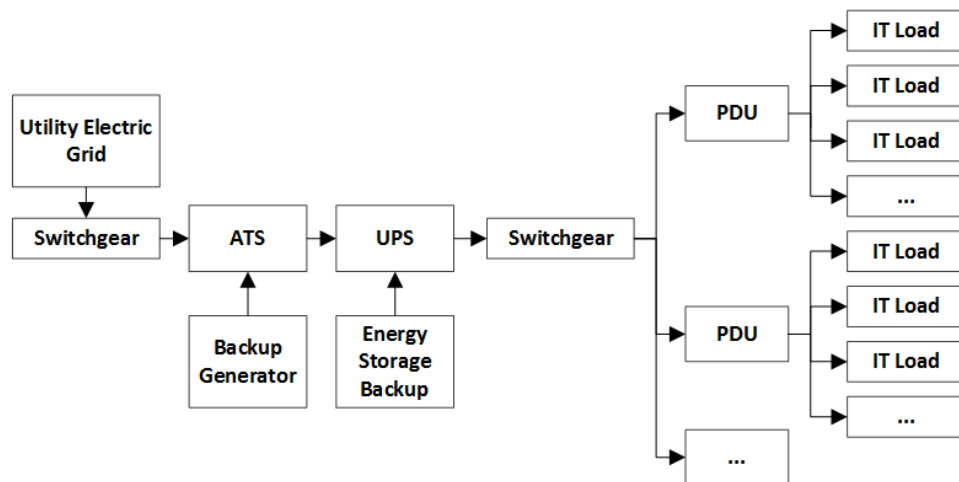


Figure 1. Schematic of the typical power distribution network in a DC.

All of this equipment produces heat, especially the servers, which needs to be removed to keep the DC equipment within its safe operating temperature range. The American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) Technical Committee 9.9 publishes guidelines for safe operating temperatures and humidity levels [16]. Figure 2 shows a basic

schematic of the path of the heat from the IT equipment to the outside environment in a typical DC. Removing the heat from the IT room itself is performed with a cooling unit such as a computer room air conditioner or computer room air handler (CRAC or CRAH), terms that are often used interchangeably but do have a distinct difference. Both consist of many separate components, such as fans, filters, heat exchangers, pumps, valves, and sometimes an integrated humidifier. They bring in the high temperature exhaust air from the DC facility and transfer the heat to a separate working fluid through a heat exchanger to carry it out of the DC. In the case of a CRAH, this working fluid is typically water and a separate chiller removes the heat from the water before it re-enters the CRAH. If it is a water-cooled chiller, a cooling tower is the final step in the process to remove the heat from the DC facility entirely. This is the process illustrated in Figure 2. In the case of a CRAC, the working fluid can be water, or an alternative refrigerant, and the heat is removed from the exhaust air using a heat exchanger coupled with a DX cooling process. The type of working fluid would dictate the last step in removing the cooling load to the outside environment, but in many cases it is simply forced convection over a cooling coil in a unit outside the building. In general, the cooling system is the biggest source of power losses and inefficiency in a DC due to the limitations of the standard refrigeration cycle. As a result, a significant amount of IR&D is focused on providing more efficient cooling.

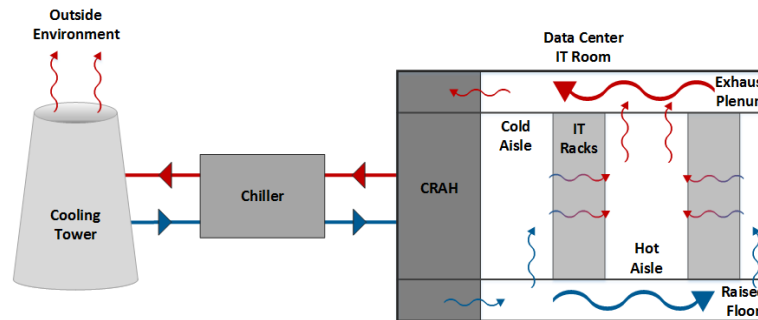


Figure 2. Schematic showing the flow of the heat out of the typical DC.

Lastly, there is an array of different pieces of miscellaneous equipment also contributing to the electrical power requirements, the cooling load, or both. Most are related to the necessity for the DC to be a hospitable environment for maintenance and IT workers. The DC lighting is part of its electrical load and needs to be included in the operating costs, and as lighting is not 100% efficient, the electrical losses need to be included in the calculation of the total DC cooling load. Computer monitors also present an additional, albeit small, source of heat and additional electrical power. Even the workers themselves need to be included in the calculation of the total DC cooling load, as the body radiates heat to its environment. Figure 3 shows a schematic of the typical DC demands and their losses due to the different components outlined above. By measuring all of these demands and losses it is possible to calculate the overall energy efficiency of the DC.




Figure 4. An example distribution of the DC power consumption for a DC with a monthly average PUE of 1.51 in August in northern California [20].

Figure 4 presents an example of the breakdown in power consumption by the major components of a DC, this one with a PUE of 1.51. As is expected, the vast majority of the DCPI consumption is for cooling. The self-reported Industry standard PUE was 1.7 as of 2014, up from 1.67 in 2013 but down from 1.8 and 1.89 in 2012 and 2013, respectively [21]. Some companies report significantly better performance: e.g., Google reported a 1.12 fleet-wide quarterly PUE for 2015Q2 and a 1.12 trailing twelve-month average, both with a 1.09 minimum and 1.24 maximum [22].

Although PUE is the standard metric used throughout Industry for reporting the energy efficiency of their DCs, it is not without its weaknesses. The first and main weakness of PUE is the lack of a rigorous standard for measuring the PUE. Although the ideal location for the IT power measurement is directly before each individual server, this can be logistically challenging and costly due to the sheer number of servers in some DCs. As a result, often the IT power measurement is performed before the PDUs or even as far back as a UPS if IT equipment is the only load. This incorporates PDU and UPS, both DCPI, inefficiencies in the measurement and must be accommodated for. The same can be said for the power measurement made of the entire DC (IT & DCPI), which can be made in a large number of locations. Compounding the issue of where to perform the measurements in the power distribution chain is the length of time a reported metric over which the measurements are averaged. As a DC's environment and IT load can vary significantly and frequently, annualized PUEs are preferable but not always the format that is reported.

An additional, frequently overlooked, weakness of using the PUE metric is that it can have the unintended consequence of dissuading DC operators from upgrading to more energy efficient servers and IT equipment. Examining the equation for PUE above, it can be seen that if the "IT Equipment Energy" increases relative to the "DCPI Energy" and the other terms in the numerator, the PUE actually converges toward 1, the ideal PUE. As a result, it is possible to install more energy efficient IT equipment and actually degrade a DC's PUE. As long as IT power or energy is in the denominator, a more energy costly server for the same IT performance will in fact improve the PUE. This is troublesome as servers are becoming more energy efficient, in terms of the amount of power required to perform the same operation, and DC operators should not be discouraged from taking advantage of those advancements [23].

To realize another weakness of the PUE metric, recall that a standard rack server typically has its own cooling fan installed. This means that for most PUE measurements the electric power used to run this fan is included in the "IT Equipment Energy" component as opposed to the "DCPI

Energy” component. A manufacturer or DC operator could feasibly increase this fan speed as much as possible, allowing the cooling unit fan speed to be lowered, and the DC PUE will decrease because the “IT Equipment Energy” component increases while the “DCPI Energy” component decreases. The DC fundamentally is not more energy efficient, but the lower PUE gives the impression it is.

One way to rectify this is to have the DC output be represented by its IT service performance as opposed to the amount of energy required by the computing hardware. The ideal DC energy efficiency metric would be to measure the number of specific IT transactions, e.g., individual e-mails, per watt-hour required to operate the DC. However, as not all DCs perform the same function, nor even only one function, developing this metric has proven difficult. The Standard Performance Evaluation Corporation has a server-class computer equipment power and performance characteristics benchmark: SPECpower_ssj[®]. This benchmark can be summarized as the amount of electrical power required by a server to simultaneously perform a given number of specific server-side java operations. Tier44, together with Underwriters Laboratories, maintains the PAR^{4®} benchmark [24]. This benchmark uses ticks-per-second (Tps) as opposed to java operations and measures the power consumption in the off, idle, and fully loaded states while also taking into consideration the peak power consumption of the server during testing. Combining one of these metrics or another similar one with the PUE ratio would better quantify the energy efficiency of the DC performing its mission: providing IT services. Otherwise any energy efficiency evaluation focused on PUE is only obtaining half of the total picture.

Metrics commonly used to examine the efficiency of the cooling system include: the Cooling Capacity Factor (CCF), the Rack Cooling Index (RCI), and the Return Temperature Index (RTI). CCF is a relatively simple way to determine whether the running cooling capacity from a traditional cooling unit (CRAC or CRAH) system matches that from the critical IT load [25], [26]. The operator divides the total of the rated capacity of the active cooling systems by 110% of the IT critical load, all in watts. The ideal CCF is 1.2, with larger ratios presenting opportunities to turn off unnecessary cooling equipment. If a DC has a large CCF ratio while also having issues removing all of the IT load heat, i.e., hot spots, then it can be concluded that the DC likely has very poor air management practices. Using CCF it is very easy to examine the design and performance of a DC’s cooling system without needing sensors or detailed information.

With RCI and RTI, it is possible to quickly generate more precise conclusions about what is wrong with a DC’s cooling system design once the necessary sensors have been installed. RCI has two different versions, high and low, both with 100% being the maximum and ideal value [27]. RCI_{HI} is the percentage of server intake temperatures below the maximum recommended value for the IT equipment. Anything less than 100% means that there are hot spots in the DC. RCI_{LO} is the percentage of server intake temperatures above the minimum recommended value, with anything less than 100% shows that there are cold spots in the DC, or areas where the cooling system is performing excess work and wasting energy. Calculating both of these metrics can be done using rack-level server inlet temperature sensors and information regarding the IT equipment manufacturer’s recommended operating conditions. RTI is more difficult to determine than RCI and identifies both bypass and recirculation in a DC [28]. Bypass is when low temperature air from the cooling unit enters the high temperature exhaust air stream without going through the IT equipment. Recirculation is when air that has already passed through the IT equipment is pulled back through or around the racks and goes through the IT equipment again. An RTI between 80-120% is the target range, with <100% implying net bypass and >100% implying net recirculation over the entire DC. RTI can be calculated from rack-level temperature

sensors, cooling unit supply and return air temperature sensors, and the cooling unit volumetric airflow rate (static value). Using a combination of both RCI and RTI, and the information that goes into computing both, it is possible to generate a very good understanding a DC's cooling system performance.

Lastly, it bears mentioning that in some locations water availability is severely restricted due to droughts or generally arid climates. Since many of the most effective DC cooling technologies use water as the working fluid because of its greater thermal transport properties compared to air, some companies track, report, and try to improve their Water Usage Effectiveness (WUE) [29], [30]. This metric can either be specified at the site level, dividing the number of liters of water used by the DC facility by the total IT energy consumed in kWh, or at the source level by also including the water use of the power plants and other facilities that provide the electric power to the DC.

1.2.3 Data Center Design Best Practices for Energy Efficiency

Multiple agencies and companies have released high level “best practices” design guides for reducing DC PUEs [31], [32], [33], [34], [35]. A good, recent example is Google, which advertises the five different best practices listed below [36]. These are just a handful chosen to give a brief overview and are not intended to be an exhaustive DC design guide.

1. Measure PUE
2. Manage airflow
3. Adjust the thermostat
4. Use free cooling
5. Optimize power distribution

Measuring the PUE seems very complex, particularly given the number of different components encapsulated in Figure 4. However, just measuring the DC electrical power demand to the utility grid and the power distributed from the PDUs to the IT equipment can provide a PUE value, although minus any information on the underlying distribution and of limited accuracy for comparison to different DCs as explained in Section 1.2.2. Also, it is key that any reported PUE value be averaged over a representative amount of time. Due to the significant day-to-day and seasonal variations in the PUE for a single DC, both in terms of IT load and environmental operating conditions, the standard practice is for an instantaneous PUE to be averaged over at least an hour and an averaged PUE to be averaged over an entire calendar year [17], [19].

Given that the second most energy-intensive operation for a typical DC after the IT systems themselves is cooling, it is no surprise that the next three “best practices” relate to cooling. First, a significant issue for air-cooled DCs is the efficiency of delivering cold air from the cooling unit to the servers to remove the heat they produce, and then back to the cooling unit to remove the heat from the IT room entirely. Managing this airflow in an effective manner can significantly reduce cooling inefficiencies and allow a DC operator to better match the produced cooling airflow to the actual cooling load. The closer the match between the two, the lower the electrical power required to dissipate a given cooling load. Reducing cooling inefficiencies also reduces the maximum cooling system design specification, decreasing both initial capital costs and operating expenses. Second, running the IT equipment at higher temperatures further reduces the amount of energy required to cool the equipment as the cooling unit output can be at a higher

temperature. The ASHRAE TC 9.9 periodically reexamines its recommended operating temperature and humidity ranges, expanding both in recent years [16]. The cooling unit input (IT room exhaust air) temperature will also increase, keeping the difference in temperature across the cooling unit relatively constant. The energy savings primarily come from the increased temperature difference between the environment outside of the DC and the exhaust air, increasing efficiencies of the chiller and the heat exchangers that transfer the heat out of the DC. Lastly, free cooling can be summarized as using the environment to dissipate heat from the DC, as opposed to solely relying on mechanical refrigeration techniques. This can be done using a few different, readily available natural sources: nearby large bodies of water, outside air, and even the Earth. Compared to mechanical refrigeration cycles, these methods require significantly less electrical power but can often only satisfy the cooling demands for a portion of the year.

To explain the last “best practice,” it should be noted that it is often misunderstood and thought that backup generators, UPSs, PDUs, and even power line cabling have constant power losses, or inefficiencies, relative to their loads (i.e., their percent overall efficiency is constant). If the UPS is sized for a 100 MW DC but is only being used for a 10 MW DC, or at 10% utilization, then its operation will be less efficient than if it were running at 80% utilization or higher. So, in addition to purchasing energy efficient power distribution equipment, the DC designer should also size it carefully or purchase modular units that can be easily tailored to the current load. This concept applies to more than just the power distribution, as well.

1.3 Technology Taxonomy

The results of the Industry technology survey showed that the technologies and methodologies used can be divided into three main categories: system efficiencies, efficient cooling, and efficient computing. Figure 5 shows the resulting technology taxonomy.

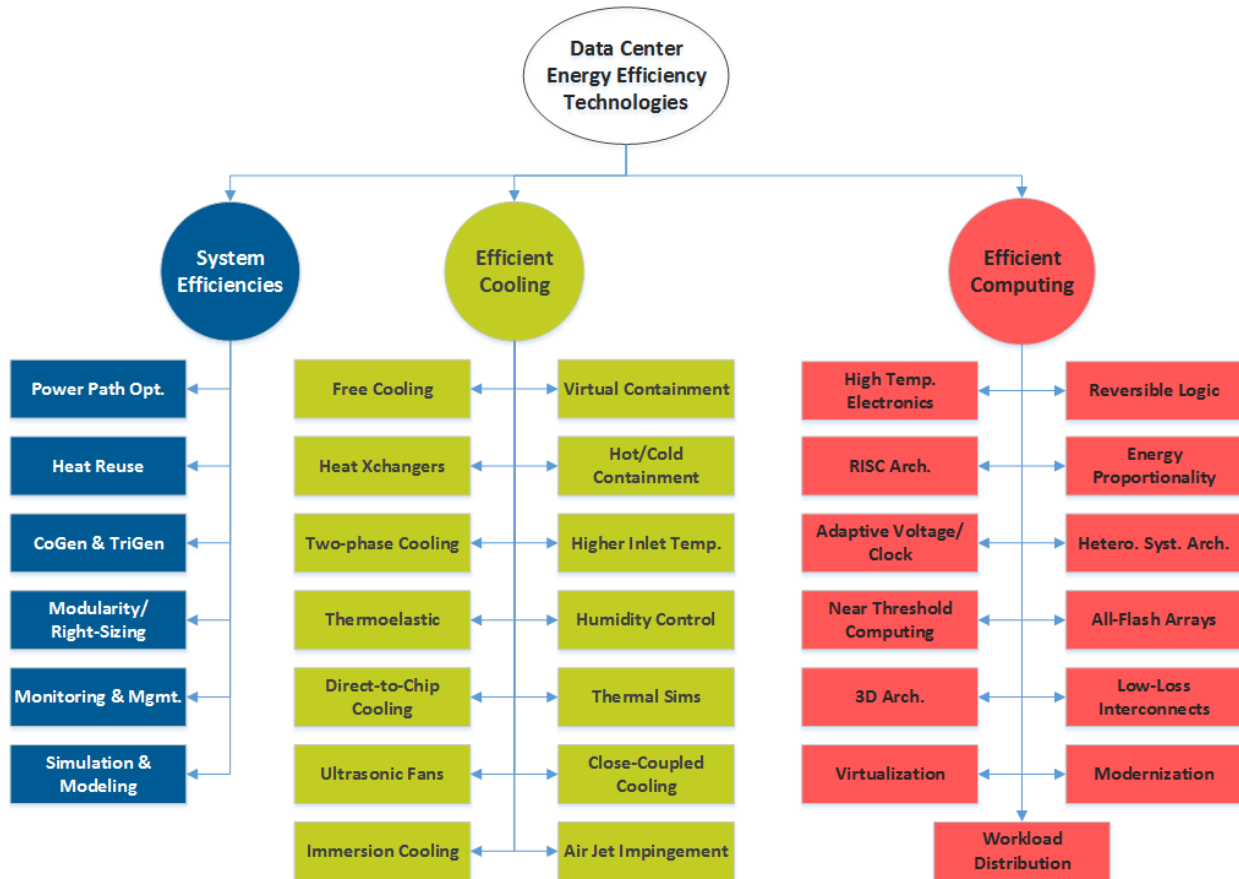


Figure 5. Technology taxonomy of data center energy efficiency technologies.

Based on the technology survey and general DoD DC needs, a number of different technologies and methodologies show significant potential for increasing DoD DC operating efficiencies and therefore mission effectiveness. These technologies can be separated into three different groups, each of which is described in more detail in the subsequent sections of this report: management methodologies (Section 2), low-capital technologies (Section 3), and high-capital technologies (Section 4).

Table 5. Promising technologies or management methodologies.

Management	Low-Capital	High-Capital
Data center floor monitoring and management	Airflow management	Direct-to-chip cooling
Right-sizing throughout data center lifetime	Economizer modes for chillers	Waste heat reuse for space heating
Decommissioning and upgrading	Power path optimization	Co-generation and tri-generation
Virtualization for flexibility	Supplemental evaporative cooling (e.g., foggers)	
	Close-coupled cooling	

2 Management Methodologies

During the review of Industry efforts to improve DC energy efficiency, a group of “best practices” surfaced that do not fall strictly into the category of a hardware technology or a hardware management software technology. Generally, these can be considered DC management techniques or methodologies that may or may not involve new technology; more importantly they involve significant decisions on how the DC is managed. The most promising DC management methodologies found during the review are: instrumenting and metering the IT and DCPI to allow for informed DC management; right-sizing the DC through its lifetime; decommissioning, upgrading, and consolidating both hardware and software systems; and virtualization of the applications, storage, and other functionalities using a software defined DC approach. The following three sections describe each methodology and its potential impact in more detail.

2.1 Data Center Floor Monitoring and Management

One of the key aspects of targeting and enacting energy efficiency improvements in a DC is sufficient knowledge of its current state and performance. For example, the typical response to servers overheating is to turn down the cooling system set point, lowering the temperature of the air reaching the servers. Knowing that the temperature of the inlet air increases more than expected between the cooling unit outlet and the server inlet would tell an operator that potentially more cost effective actions can be taken to prevent overheating by making the cooling system more efficient instead of simply turning up its output. With the goal of providing DC operators with more actionable information on the state of the DC, the EPA includes a real-time server data reporting requirement as part of its ENERGY STAR[®] computer server certification [37]. All certified servers can provide input power consumption, inlet air temperature, and CPU utilization information to non-proprietary software through the server’s management port. Some servers are providing even more data, such as the level of airflow in the server and its temperature increase upon exiting through the exhaust [38]. Unfortunately, there is no single standard for the formatting of the server data. In a DC with a wide variety and range of server models and IT equipment, collecting the information directly from the servers can be challenging to implement in a cost-effective manner and might still produce an incomplete picture of the DC. Alternatively, many vendors offer stand-alone power, temperature, humidity, and airflow sensors to allow a DC operator to monitor the environment within their DC [39], [40], [41]. Many vendors also offer PDUs or entire power distribution systems with built-in power metering capabilities [42], [43], [44], [45], [46]. All of this information can be used to track surges and peaks plus perform power usage, server utilization, or thermal audits. Figure 6 illustrates an example layout for how this information can be used to display DC power, thermal, and utilization data. Empowered with an intuitive and accurate map of the DC, an operator can quickly identify hot or cold spots and chronically idle machines, among other things.



Figure 6. An example display for how server temperature and power information can be used to monitor the state and performance of a DC [47].

A large number of vendors market software products to help collect, analyze, and visualize this information [48], [49], [50], [51], [47], [52], [53], [54], [55]. Typically referred to as Data Center Infrastructure Management (DCIM) software, its reported benefits beyond the generation of displays like the one shown in Figure 6 include the ability to forecast the effects of changes to the DC with built-in models, the ability to reduce costs by optimizing power consumption through informed load distribution and allowing the operator to turn off unnecessary hardware, and the ability to reduce downtime from critical events with advanced warning and faster troubleshooting [56]. DCIM software can often collect data from both the server themselves and standalone sensors and meters. One vendor touts the ability of their DCIM software to boost DC operating and energy efficiencies by 10-30% [57]. Although the use of fully-integrated DCIM software with a DC monitored by a full suite of sensors is the ideal path to optimize performance, minimize energy requirements, and identify opportunities for improving energy efficiency, for many DCs it may not be cost effective. In fact, many in the Industry are currently disillusioned by lofty promises and lacking execution, combined with ballooning costs and implementation timelines [58]. A partial solution of sensors combined with smaller scale or even custom data analysis software is still an effective path to a more energy efficient DC, and, at this point, may be the most cost-effective route.

Another tool operators can use to better understand the DC IT room environment and predict the impact of potential changes is a computational fluid dynamics (CFD) simulation. Multiple vendors offer simulation packages specifically designed for traditional DCs, and they can be valuable in understanding the expected impact of the airflow management or containment strategies discussed in Section 3.1 [59], [60], [61], [62]. Combined with data points from distributed sensors, they can be used as a model of the typical airflow in the specific DC. It is important, however, to understand the limitations of CFD. Without real sensor data from the DC, the simulation cannot be expected to accurately estimate the current state of a specific DC nor the exact result of installing airflow management improvements. Even with real sensor data, the results could still have a large degree of error. It is also essential to realize that any single simulation result only represents a single point in time for a DC whose IT demands and environmental conditions vary widely over the course of a year, over the course of an individual day, or even over the course of an hour. Performing a representative, annualized CFD simulation would be computationally prohibitive and still carry accuracy concerns if used to calculate quantitative metrics.

2.2 Right-Sizing throughout Data Center Lifetime

When designing a DC, it is very common to predict what the maximum IT load will be and design the initial DC around this number. Right-sizing is the practice of initially installing only the necessary capacity and adding more as it is needed. Figure 7 shows the amount of capacity “waste” resulting from (a) installing 100% of the eventual DCPI design capacity at the commissioning of the DC compared against the “waste” when (b) right-sizing the DCPI throughout the DC lifetime with two supplemental installations. The actual load (green curve) and the expected load (dashed gray curve) shown in the plots were obtained from the actual installations of some Schneider Electric customers and surveys of others [63]. The data summarized in Figure 7 represents the average and Figure 8 shows the distribution, which follows a bell curve that leans toward lower utilization levels. Combined, these figures illustrate that, on average, being flexible and expanding to meet new IT demands as they come is a significantly more energy and cost efficient approach compared to trying to deliver 100% of the envisioned eventual IT demand at the start.




Figure 7. Designed DCPI capacity and actual load over the lifetime of a DC using a traditional (a) design methodology and (b) by right-sizing through the DC lifetime [63].




Figure 8. The distribution of the ultimate DC utilization levels for a typical DC [63].

There are a couple of different methods being used in Industry to right-size the DC. The most popular is the use of containerized prefabricated DCs or DCPI to simplify the process of adding new capacity. These typically come in International Organization for Standardization (ISO) shipping containers, and there are numerous vendors in the market [64], [65], [66], [67], [68], [69], [70], [71]. Some vendors have expanded to even smaller self-contained DC systems, with models offered at the individual 42U rack size and smaller [72], [73], [74], [75]. The ruggedized versions of these are particularly attractive for the DoD unique needs mentioned in Section 1.1.

The global market for modular, prefabricated DCs is projected to have a 31.9% annual growth rate through 2019, reaching \$26 billion from an estimated \$6.5 billion in 2014 [76].

Prefabricated, modular systems have significant benefits, particularly in terms of deployment speed (60% reduction [77]) and reduced overall capital costs from avoiding purchasing ultimately unnecessary capital equipment (27.2% possible reduction based on one modeled case [78]). One drawback for these systems relates to the site infrastructure requirements for integrating the prefabricated system, something which could also easily change from generation to generation [77]. Another is that a prefabricated, modular ISO shipping container filled with racks of servers will have a higher initial capital cost than the equivalent number of racks and servers purchased individually. This is because the manufacturer of a modular system adds on the cost to assemble and configure it in their warehouse.

Installing IT capacity over time can also be accomplished through the build-up of the IT racks and DCPI equipment section-by-section inside of the overall DC facility [70], [79]. Taking the modular components out of shipping containers significantly increases their ability to adapt to the DC facility's infrastructure, even if it was designed for a previous generation. Section-by-section modularity inside the DC facility is most effective when paired with a close-coupled cooling technology, which will be discussed in Section 3.5, because otherwise the room-level cooling system could be exerting extra effort to cool a vacant section of the DC that hasn't yet had IT equipment installed.

There are two ways over-sizing produces excess costs for a given IT load: initial capital expenses and operating expenses. A DC's predicted IT load and the overall installed capacity can be used in conjunction with the PUE to predict the electrical power portion of its operating expenses. Figure 9 shows that a DC's PUE improves at higher utilization levels. As the curve makes clear, a given IT load will require more electrical energy at lower DC utilization levels. Schneider Electric has published a number of whitepapers predicting a potential 50% reduction in electricity losses for one modeled system [80], a 10-30% potential reduction in power consumption on average [31], and a 32% operational cost savings over a 10 year lifecycle for another system [78] by right-sizing the DCPI over the DC lifetime.




Figure 9. The PUE of a typical DC as a function of the utilization level, based on a detailed DC model [80].

Right-sizing the DC over its lifetime also has a positive impact on the initial, up-front costs for a new DC, which are typically capital expenses. The Total Cost of Ownership (TCO) model for one specific use case predicted 27% capital costs savings over the 10 year DC lifecycle [78]. The effect is significant for the DCPI, with a model predicting the potential to eliminate up to 50% of the DCPI TCO over the 10 year DC lifecycle [31]. Figure 10 provides a breakdown of the specific sources of savings predicted by the model for both the (a) capital costs and the (b)

operational costs mentioned earlier in this section. The most significant contributor to both the 10-year capital and operational cost savings estimates comes from avoiding overbuilt capacity.




Figure 10. TCO breakdowns over (a) major capital expenses and (b) major operating expenses for a modular DC using prefabricated containers [78].

All of the savings estimates presented thus far rely on several assumptions, which are key design issues for modular DCs. The first is that there appears to be no significant cost associated with any potential DC downtime during the installation of new IT capacity, or there is no downtime at all. For some applications downtime of any length cannot be tolerated. Even when it can be, there are likely associated costs, which will vary from case to case. The second is that, at a high level, the DCPI has been sized to accommodate additional components that would be added over time or has a modular design itself. For example, some DCs require completely new utility grid substations, ATS, and switchgear that would be extremely costly to upgrade only a couple of years after its installation. Another example is that many DC cooling systems consume significant amounts of water. As the IT capacity increases, more heat needs to be removed and therefore more cooling water is consumed (e.g., evaporated into the outside air from a cooling tower), requiring the local water source to be able to accommodate the increases in demand. The last main assumption is that the DC has the ability to closely monitor the current state of the DC (e.g., overall IT capacity and DCPI utilization) to accurately forecast the timing of future infrastructure needs to allow enough lead time.

2.3 Decommissioning and Upgrading

Decommissioning idle servers is one way to gain some of the benefits of a right-sized DC for an oversized existing facility. Often servers end up sitting idle for significant lengths of time because they essentially slip through the cracks in the normal day-to-day of running the facility. Simply cataloging the equipment that is no longer necessary and taking it out of the DC can result in increased energy efficiency, with one study estimating \$500 in annual energy savings per 1U rack server decommissioned [35]. Aside from energy savings, decommissioning equipment can also save on software licenses and provide a one-time influx in capital from recycling [81], [82].

Once unnecessary equipment has been decommissioned, energy efficiency improvements can still be gained by modernizing the equipment that remains. Over the past decade servers have become more efficient, with the average computational performance compared to the required IT

power curve (as a function of utilization) increasing by approximately 4x and the server power consumption while idle also decreasing significantly [23]. Regardless of the current trend of increasing individual server maximum performance capabilities, a newer server will require less electric power for a given computational load. Also, particularly old equipment may not have the power management techniques present in modern systems [31]. For example, a modern server has an idle, or no-load, power consumption that is significantly less than its power consumption at 100% utilization. Some older systems have identical power consumption at idle and at 100% utilization, remaining constant no matter the utilization level of the server and leading to calls for energy-proportional computing [83]. This drastically increases the server electrical power consumption for a given computational load compared to the standard systems readily available today. A test case conducted in collaboration with the EPA showed that replacing older servers with EPA ENERGY STAR[®] certified servers resulted in a 54% drop in power consumption for those units, with certified servers being 30% more energy efficient on average than a standard unit [35], [84]. Another test case had a DC that consumed a maximum of 10 MW (including DCPI, not just IT) replace its 2U servers from 2008 that delivered 133k Tps/watt with 1U servers from 2010 delivering 700k Tps/watt. This resulted in an 81% reduction in power consumption for the same application load and increased the facilities maximum potential computational capacity by more than 9x [24].

There are two technologies that bear mentioning here as they represent the new generation of IT equipment: blade servers and flash data storage. Contrary to the standard “xU” server configuration, a blade server contains only the main computing electronics and outsources the majority of the power conditioning, localized cooling fans, and data connections to a vendor-specific chassis that holds the blades and sits in the 42U rack. From a pure IT performance standpoint, blade servers allow for significantly higher levels of DC IT capability within a smaller footprint. Although blade servers concentrate the heat, causing higher exhaust temperatures relative to the inlet, there are some advantages to this for DC energy efficiency [85]. As blade servers have a smaller volume than a comparable (in terms of kW IT power) “xU” server, the cooling load is at a higher temperature, sometimes delivering up to twice the increase in temperature between the inlet and outlet cooling air. Not only does this increase the opportunities for waste heat reuse, as the air can exit in the range of 115-125 °F, but less air flow is required to cool a given IT load and the cooling unit is overall more efficient. Less air flow means not only less electric power consumed by fans, but also smaller fans with lower capital costs and more flexibility to match the current cooling load for efficient operation. The main drawback has already been mentioned: the need to use vendor-specific blade systems and chassis as there is not the same level of standardization across vendors as with standard “xU” servers. Also, existing cooling units for room-based air-cooled DCs might not be able to accommodate the higher, localized cooling loads. Even if they can be accommodated in the DC, the cooling units likely would not be able to accommodate consolidating them to make room for new IT capacity. The cooling systems often need to be upgraded as well to fully take advantage of the performance and energy efficiency of blade servers.

Although the largest contributor to the IT energy load are servers, data storage can be an appreciable component. A Forrester Research study conducted at the request of Pure Storage, a flash-based data storage vendor, showed significant benefits when using all-flash arrays (a type of SSD) instead of hard disks [86]. For a composite, representative DC the all-flash arrays had a 102% ROI and 14-month pay-back period over the first three years of their use. The storage cost approximately \$947,000 to implement and operate, representing nearly \$86,000 in power and cooling savings over those three years. These power savings come from a significantly lower

cooling load, higher maximum safe operating temperatures, and significantly faster ramp-up and ramp-down times. Unfortunately, SSDs in general are significantly more expensive per byte of storage than hard disks. Hybrid solutions, or even individual hybrid drives [87], that combine both SSDs and Hard Disks are a cost-effective way to take advantage of the benefits of SSDs. As this section shows, even just keeping a DC's IT hardware and applications up-to-date can significantly improve their energy efficiency, although even this can prove costly.

2.4 Virtualization for Flexibility

At its base, virtualization is the disassociation of software applications from specific physical servers. Traditionally, an application is hosted on a single server, shown as the Hardware-Defined Data Center (HDDC) in Figure 11. When that application is not in use, the server sits idle. The application is limited to the processing capabilities of that specific server, as well. Virtualized applications, however, are free to operate on a generalized computational capability built across many different specific pieces of hardware, as shown by the two configurations on the right-hand side of Figure 11. Servers are not the only IT system that can be virtualized. Virtualized data storage, often referred to as cloud storage, operates by the same principles as storing or accessing a file does not necessitate designating specific storage drives as the file can be shared between physical locations. Virtual machines are another example, where user desktops are not associated with a specific, attached computer or even storage drive and can be accessed anywhere on the network. Two of the major providers of DC virtualization software are VMware [88] and Microsoft [89].

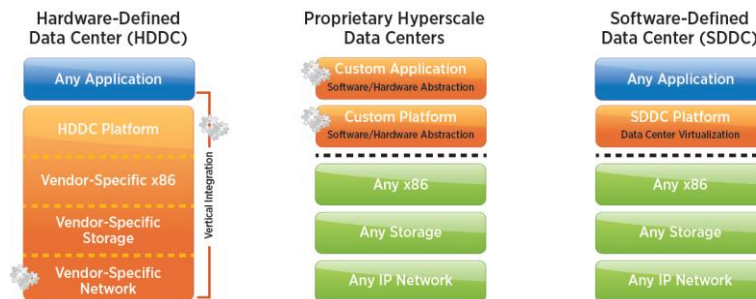


Figure 11. Comparison between traditional Hardware-Defined Data Center (HDDC), proprietary virtualized DCs, and the VMware Software-Defined Data Center (SDDC) [90].

The primary energy efficiency gains from virtualization are tied to the effects discussed in Section 2.1 for right-sizing the DC. For example, consider a DC with 100 servers for 50 different applications, a common 2N redundancy configuration. If one server is dedicated to an individual application, with a second server functioning as a backup in case the primary server fails, then two full servers could be sitting idle and drawing power unnecessarily when the application is not in use. Even when the application is in use, it commonly does not use 100% of the main server, so less than 50% of the IT capacity dedicated to that application is being used. This is a significant contributing factor to the very low utilization levels shown in Figure 8.

Virtualizing the applications can increase the overall utilization factor by as much as 80% [88]. This is due, first and foremost, to the ability to host multiple applications on a single server and split larger applications between multiple servers. If applications can be quickly and easily transferred to a group of servers in the case of a failure, what was previously an entire backup server can be replaced with 20% capacity here, 30% capacity there, and so on. Now, a server

running at <100% utilization can theoretically serve as a portion of the backup if another server fails and idling servers are no longer necessary to provide the desired level of IT availability. The above description applies to data storage, as well, as previously individual hard drives were backed-up through mirroring the drive to a second.

While this approach will not reach the level of availability provided by a 2N configuration, it can come extremely close at a significantly lower overall cost. For example, one analysis predicts a 10-40% potential reduction in DC power consumption with server virtualization as each server that is eliminated can save approximately \$380 per year in electricity costs [31]. VMware predicts that its entire SDDC software package can reduce operating expenses by as much as 56% [91]. An outside study of their product attributed it for a 50% reduction in total IT costs [92]. The VMware website also predicts an 80% maximum reduction in energy costs, a 70% reduction in new application provisioning time, a 20-30% lower cost per application, and a 15:1 reduction in the necessary hardware [88]. The VMware Distributed Power Management (DPM) software, one option among their suite of software tools, likely contributes heavily to these energy savings as it has the ability to dynamically power on and off specific servers as the IT workload varies [93]. Put another way, an existing traditional DC switched to virtualization can significantly decrease operating expenses by reducing the number of necessary and active physical servers while only sacrificing a relatively small amount of IT availability.

At this point it is important to realize the limitations of virtualization with respect to improving DC energy efficiency. As has been described, virtualization allows a DC operator to consolidate the IT load and eliminate a large number of physical servers from the racks. A major, but not yet mentioned, assumption is that the DCPI is able to accommodate the server consolidation. If a DC uses a CRAC or CRAH cooling unit for room-based air cooling, it might not be able to remove the higher localized cooling loads produced by servers operating at higher utilization levels. Suppose 100 servers spread throughout 10 different racks operating at 10% utilization are replaced by 13 servers operating at approximately 80% utilization, 10 in racks by themselves with the remaining 3 spread out. The cooling unit might not be able to support consolidating the 13 servers into only 1 or 2 racks or support the addition of more servers to take advantage of the vacant rack space. If that's the case, the energy efficiency benefits could be minimal compared to the cost of transitioning to a virtualized DC to get down to only those 13 servers. Supplemental, localized cooling could be used, but that would increase the capital and operating expenses, negating some of the gains resulting from eliminating physical servers. Therefore, virtualization is often performed in conjunction with other upgrades, such as to close-coupled cooling, and in many cases cannot be thought of as a stand-alone opportunity for improving the energy efficiency of a DC.

2.5 Summary and Recommendations

This section focused on the specific DC management methodologies and techniques that Industry has demonstrated to significantly increase the energy efficiency of their DC operations. The first and foremost step recommended for any DC is performing an exhaustive inventory and at the least decommissioning unnecessary and legacy systems. DC IT technology has come a long way in terms of its energy efficiency even just over the past decade and using obsolete equipment significantly increases DC energy consumption. Requiring traditional DoD IT vendors to report Industry standard energy efficiency performance metrics for their equipment and to meet specific goals, such as requiring ENERGY STAR[®] certification, would go a long way to ensuring a DC

is energy efficient. The next necessary step is ensuring that sufficient and appropriate data is collected on the current state and performance of the DC. Whether this takes the form of a fully-integrated DCIM solution or a smaller, custom solution, continuous access to power, thermal, and utilization information throughout a DC will drive energy efficiency optimization and identify infrastructure improvement opportunities. This information can also potentially be paired with CFD models of the DC floor airflow. Another step any DC operator can take is ensuring that their facility is appropriately sized throughout its lifetime. The tendency to size the facility and install DCPI according to the projected DC needs at the end of its lifetime, typically as much as 10 years, results in significant waste as initial estimates have been 50% higher than the actual final need, on average [63]. Lastly, when an organization is preparing to perform a major overhaul of their DC, virtualization of the applications and data storage increases the overall density and energy efficiency of a DC but often requires concurrent upgrades to the DCPI. But once implemented, specific tools that automatically power on and off individual servers based on current IT demands can help minimize the power consumed by unnecessary, idle servers.

The key limiting factor for the DoD to leverage these techniques, however, is that the DoD is not managed in the same way that most companies are managed. As mentioned earlier, the DoD focuses on providing the highest level of mission effectiveness within a given budget as opposed to reducing the cost to provide a given service and thereby maximizing profits. Within the DoD, decreasing the cost of something, such as the cost to run a DC, can mean a reduction in that organization's overall budget and the reallocation of those funds elsewhere. Therefore energy efficiency improvements are often seen as leading to negative impacts on the organization that operates the DC.

Another significant limiting factor for the DoD comes from the nature of the DC itself, and many large companies are faced with a similar issue. A DC is seen as an IT capability, and is managed by IT organizations. However, that organization rarely, if ever, sees or pays the bill for the electricity consumed by the DC. That is handled by the organization managing the facility the DC is in. As a result, one organization is often in charge of purchasing and managing the networking equipment and servers while a completely separate organization is in charge of purchasing and managing much of the DCPI as well as paying the utility bill (a.k.a., stovepiping). The Air Force DC Consolidation Energy Efficiency Pilot at Edwards Air Force Base illustrates an approach to overcoming this limitation [94]. It makes use of an Energy Savings Performance Contract (ESPC), a contracting vehicle already in use by the Department of Energy (DOE) [95]. When an ESPC is awarded, the company is able to come into an existing DC and completely refurbish the facility using its own funds and is given management responsibility for the duration of the contract, typically 10 years. The DoD pays the company a flat rate throughout the length of the contract, such as the DC's original operations and maintenance (O&M) budget. The company is then able to pocket, as profit, the difference between what it costs them to refurbish and run the DC and what the DoD previously paid in O&M. From the DoD perspective it does not have to put forth the upfront capital costs for upgrading the facility and at the end of the contract period it gains ownership of the capital equipment for the rest of its operating life, sometimes an additional 5 years in the case of some DCPI. This type of contract does come with risks, however. The vendors make the proposals based on the information they are given about the DC, which is typically significantly less than would be provided by the management technologies discussed in Section 2.1. Therefore, the awarded vendor's business model could be very inaccurate as many DoD DCs lack the

performance information to properly inform vendor business models and the technical expertise on staff to validate proposals during the evaluation process.

These management methodologies and their associated technologies present potential advantages for some DoD unique applications, as well. As illustrated by the aforementioned Northrop Grumman CPP RWS, some DoD DCs need to be highly mobile and take little-to-no time to set up in the field. Assuming they can meet the requisite survivability specifications, prefabricated, containerized DCs might be ideal for providing immediate C4 capabilities to Forward Operating Bases (FOBs) or some larger Command Posts (CPs). Also, virtualizing the fielded DCs will significantly decrease their footprint while increasing resiliency to component failure, lowering the mass and volume of equipment transported and decreasing the associated logistics.

Virtualization can also help decrease the installation and operating costs of DCs handling Impact Level 5 and 6 Information. A more compact DC will most certainly decrease the cost of the required access controls and physically separating it from other IT systems.

3 Low-Capital Technologies

Among the array of technologies used by Industry to increase DC energy efficiency is a group that is very conducive to upgrading and updating existing DCs due to the relative ease of implementation and low cost. This section describes the low-capital technologies that stood out as potentially providing significant benefits: airflow management and hot/cold containment for air cooling, economizer modes for chillers, power path optimization, evaporative cooling using foggers, and close-coupled cooling.

3.1 Airflow Management

One of the main strategies for increasing the energy efficiency of a DC is to ensure that the task of removing heat from the IT systems is as efficient as possible by isolating the inlet and outlet cooling fluids. This is relatively simple in liquid-cooled DCs as the fluid travels in separate pipes. However, the vast majority of DCs use air as the fluid, making this task more complex. Figure 12 shows the isolation of the cold inlet air from the hot exhaust air using a raised floor and an overhead plenum in a hot-aisle containment configuration. In addition to the raised floor to deliver the inlet air and the overhead plenum to funnel the exhaust air, there is a barrier inside the main room keeping the exhaust air in the hot aisle from mixing with the air in the rest of the DC, including the inlet air in the cold aisle. To achieve hot and cold aisles in a DC, the racks must be arranged to have the IT equipment facing each other, i.e., front-to-front in the cold aisle and back-to-back in the hot aisle. Cold-aisle containment, isolating the inlet air from the rest of the DC, is also possible and the choice between hot- or cold-aisle containment should be made on a case-by-case basis. It should be noted that raised floors and separated overhead plenums are not strictly necessary or always implemented. Natural convection currents in combination with high ceilings can help create the same effects, although to a lesser degree. The airflow strategy can also eliminate the need for a raised floor, for example, by having the cooling units release the treated air directly into the open DC area and completely containing the hot aisle and exhaust return.




Figure 12. Hot-aisle containment of DC with a raised floor and exhaust plenum [96].

Hot- and cold-aisle containment systems are designed to nearly eliminate two sources of inefficiency: bypass and recirculation. The target result is a consistent cooling performance across the entire DC with no hot or cold spots, thereby allowing the cooling system to be turned down to conserve energy without endangering any equipment. Referring to Figure 12, bypass could occur when the cold inlet air being pumped into the volume underneath the raised floor reaches the exhaust plenum without passing through the server cabinets. Treating the DC room as a thermodynamic system, it strives to be at a thermal equilibrium and the lower temperature inlet air is pushed toward the higher temperature exhaust air along the path of least resistance. Without proper physical separation the path of least resistance might not be through the servers, as intended. For example, the floor underneath the hot aisle might not be properly sealed or the cold air can travel up and over a server cabinet that lacks a physical barrier above it. Recirculation is when the inlet air interacts with the exhaust air before it reaches the servers, raising the temperature of the inlet air and requiring more work from the chillers to achieve the desired air temperature at the front of each server. One cause of recirculation is insufficient inlet air flow, drawing exhaust air back around to travel through the racks again.

Significantly reducing or eliminating bypass and recirculation can have drastic effects on DC performance. One article has said that 60% of the air supply can be lost in a DC with very poor airflow management [97]. Intel reported that implementing hot/cold containment allowed for a 15 to 40 kW jump in the per-rack computing power their cooling system could handle [98]. Hot/cold air containment has been estimated to reduce energy expenses 5-10% by reducing fan energy requirements 20-25% (provided variable frequency drive (VFD) fans are installed and their output can be decreased without significant fan efficiency losses [99]) and chiller power consumption 20% [35]. When used in conjunction with a wet-side economizer (see Section 3.2.2), hot/cold air isolation techniques increased the necessary inlet air temperature by 15 – 20 °F [85]. This increased the portion of the year when the economizer could be used and also increased the efficiency of the chiller when it was needed (0.55 to 0.44 kW/ton to produce 55 °F chiller water compared to 42 °F chiller water). These results confirm that significant gains can be made using hot/cold containment to better manage the airflow in air-cooled DCs. This is a very common issue for many DCs, as well. Figure 13 shows that the average CCF seen for customers of Upsite Technologies is more than 3x the ideal value. As CCF only takes into account actively running cooling units, it is unlikely that initial DCPI over-sizing is the cause of the cooling system inefficiencies and that poor airflow management and containment are the primary causes. Better airflow management practices also produce more consistent temperatures across the DC and allow the CCF ratio to be reduced safely, resulting in significant energy savings.




Figure 13. The cooling capacity factor (CCF) values for 45 different DCs examined by Upsite Technologies [26].

3.1.1 Virtual Containment

There are multiple methods being used by Industry for air flow management and containment. Basic airflow management techniques, sometimes referred to as ‘virtual containment,’ should be implemented before any of the hot/cold containment technologies discussed below are considered [100], [101], [26]. In all DCs, even those without containment strategies, the rows should be arranged in hot and cold aisles. This minimizes the mixing of hot and cold air as all of the air in a single aisle is ideally of the same type, either inlet or exhaust. When a raised floor is used to deliver the cold air, the assumption is that it only enters the cold aisles through perforated tiles. However, many DCs have unsealed openings in the raised floor for cables to pass through or perforated tiles in the wrong location, increasing air bypass inefficiencies. Sealing these cable openings with grommets [102], [103] and ensuring perforated tiles are only in the cold aisles, not in the hot aisles or in the unused areas of the DC, will significantly reduce bypass inefficiencies. Some older DC racks have glass doors on the front with access controls, however this severely inhibits the flow of the inlet air into the front of the rack. These legacy doors can be replaced with perforated doors, available with the same access controls if necessary, to improve airflow through the rack. In most DCs there are empty spaces in each of the racks [104]. Putting in blanking panels to prevent air from flowing through the racks around, as opposed to through, the IT equipment can significantly improve cooling efficiency. Adding a 12 inch tall blanking panel to the middle of a rack can yield 1-2% in energy savings for that rack [35]. Depending on the configuration there can also be empty spaces between the racks themselves that should be sealed [105]. The next step is to completely seal off empty, unused racks from both the cold aisle and raised floor. Contact paper or other low-cost means of preventing undesirable airflow are options when the sheer number of blanking panels and other equipment needed becomes cost prohibitive. In addition to sealing off empty racks, partially blocking off the portions of the raised floor space in unused areas of the DC with removable panels reduces the chance the inlet air will go off course, reducing the unavoidable increase in temperature as it flows to the rack from the cooling unit [106].

Once all of these air flow management improvements have been implemented, energy savings will only be seen if the cooling system is adjusted to account for the new level of efficiency. Figure 14 illustrates what will happen in a typical raised floor DC if airflow management techniques are implemented but the cooling system is not turned down to take advantage of the new efficiency: wasted cooling effort.

The energy savings from implementing basic airflow management techniques can be significant. In one test case, an 85,000 square foot facility, 466 square feet of raised floor penetrations were sealed, 6,700 blanking plates were installed in racks, and the underfloor volume was reduced by 25% using a sub-floor barrier system [100]. These changes reduced the square footage of the cooled environment by 25%, increased the sub-floor static pressure by 93%, increased the sub-floor airflow rate by 49%, decreased the sub-floor air temperature by 2.7 °F, increased the cooling unit return air temperature by 3.2 °F, and allowed 29 of 60 cooling units to be turned off in the facility. These changes would save an estimated 1.6 million kWh/year, resulting in a 4 month ROI for the upgrade.




Figure 14. The airflow effects in a typical DC from implementing proper management techniques without turning down the cooling system to accommodate the increased cooling efficiency [107].

3.1.2 Hot/Cold Containment

If basic air flow management steps are not sufficient, many different containment technologies are available. Figure 15 is an example of a vertical exhaust duct, or chimney [108]. On the rack-level these vertical ducts transport the rack exhaust air to the overhead plenum [109]. A similar, row-based system stations the vertical ducts at the end of each row, connected to the hot aisle [110]. Another route is to extend walls around the hot aisles all the way up to the exhaust plenum [70].




Figure 15. Rack-level exhaust air chimney system [111].

Although requiring a little more capital investment due to its specific equipment requirements, vertical cooling technologies represent an innovative approach to hot/cold containment [112]. The technique removes the localized server and blade chassis fans and uses vertically-aligned blade servers in completely sealed racks. As shown in Figure 16, the lower temperature air in the raised floor is drawn up through the sealed racks and past the blade servers until it reaches the exhaust plenum at the top of the room. Fans at the bottom of each rack are not always needed because the air will naturally rise up through the racks as it is heated to a higher temperature relative to the air coming in from the raised floor. This technology, in effect, switches the IT equipment from front-to-back cooling to bottom-to-top cooling. This has the advantage of requiring significantly less fan power to maintain the necessary air currents; however, it also converts the cooling of individual servers from a mostly parallel configuration to a completely series configuration. The equipment at the top of the rack will always receive significantly higher temperature cooling air because it has been heated by the equipment at the bottom. As a result, there are trade-offs to consider between the reduced fan power and the conversion to a series configuration cooling loop.




Figure 16. Vertical cooling configuration [112].

While the hot/cold containment products discussed thus far do not require as large of a capital investment as some of the technologies discussed in Section 4, there are alternatives for the even

more cost-conscious. Curtains are a very common low-cost alternative solution for partial containment [113], [114], [115]. Often the curtains are made of fire-resistant materials and are designed to drop down in the case of a fire to accommodate building fire codes. Multiple companies also advertise modular containment systems, with one example of a partial containment product shown in Figure 17, that can be easily installed onto existing racks and accommodate reconfigurations [116], [117]. A comparative computational fluid dynamics (CFD) study of the example partial containment system shown in Figure 17 and a “legacy” containment system showed similar performance results (\$32,000 vs. \$35,000 annual energy savings) but with a significantly lower cost to deploy (\$36,000 vs. \$71,000) and therefore a faster ROI [118].




Figure 17. A magnetically attached modular containment system [119].

All of the hot/cold containment methods mentioned thus far promise significant energy savings and involve adding chimneys, curtains, walls, etc. onto the racks or rows to reduce or eliminate air mixing and bypass. This is a relatively simple step that can be done to increase the efficiency of a DCs cooling. However, as with basic air flow management steps, the actual energy efficiency of the cooling system does not increase until it is properly adjusted to match the new level of cooling efficiency. Figure 18, when viewed with Figure 14, illustrates that implementing basic airflow management techniques followed by cold-aisle containment, without also turning down the cooling unit output, will likely result in wasteful cold spots in the DC from lower than necessary operating temperatures.




Figure 18. A schematic illustrating the effects of cold-aisle containment without concurrently decreasing the cooling unit set points to accommodate for the new level of cooling efficiency [120].

3.2 Economizer Modes for Chillers

Every DC will require some level of heat removal, so in addition to cooling representing the second-largest electric power consumption in the average DC, it is also unavoidable. However, one way to reduce the cost of cooling a DC is to do some of it for free. Referred to as Free Cooling, some DCs use their surrounding environments as a thermal mass to transfer some or even all of the produced heat to. To facilitate this, companies have decided to locate their DCs in colder climates, such as Google's facility in Hamina, Finland [121], [122]. The Hamina DC draws in cold seawater to remove heat by passing it through heat exchangers connected to the DC cooling loop, and then recombining it with seawater to lower its temperature back down before returning it to the sea. Although it is not always possible to locate a DC in a cold climate, chiller economizer modes look to allow for Free Cooling in less optimal conditions. For example, TelecityGroup announced a 6.4 MW DC in Paris that uses free cooling and other energy efficiency improvements for a predicted 28 GWh in annual electricity savings [123]. A study of the impact of Free Cooling and other energy efficiency improvements on the National Oceanic and Atmosphere Administration (NOAA) DC, the Environmental Security Computing Center in Fairmont, West Virginia, showed a potential reduction in annual energy usage of 2.1 GWh for the 6 MW capacity facility [124]. These predictions can vary significantly based on geographic location and the specific DC, however. Figure 19 and Figure 20 show conservative estimates for the number of hours per year dry- and wet-side economizer modes could operate.

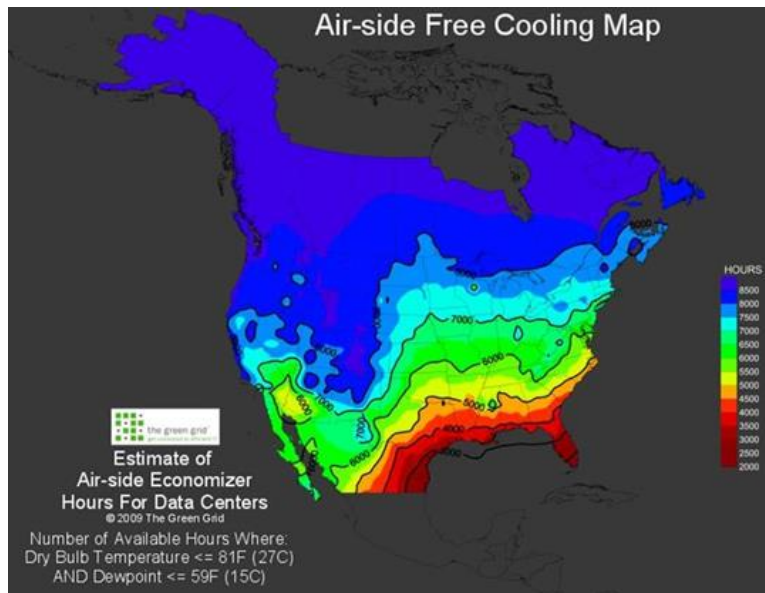


Figure 19. An example estimate of the potential number of hours per year a dry-side economizer mode could operate using a conservative performance assumption [125].




Figure 20. An example estimate of the number of hours per year a wet-side economizer mode could operate using a conservative performance assumption [125].

Economizer modes can be broken down into two main methods: dry side and wet side, based on whether the air or the water going into the cooling unit is pre-cooled [85].¹ They can be configured to run as a complete substitute for the chiller in an either/or configuration, or

¹ Note that this is not a completely orthogonal distinction as there is some contradiction between authors on the characteristics of each type of system, as will be noted in the following sections.

decoupled from the chiller such that both can operate independently and therefore simultaneously. In general, the more different modes of operation available the more precisely the cooling system can be configured for optimal energy efficiency given the conditions at the time. This is exemplified by the six different cooling options available with the Switch TSC 1000 ROTOFLY system [70]. The next two sections discuss dry- and wet-side economizer modes in greater detail, but studies by Schneider Electric have predicted that economizer modes can result in a 4-15% decrease in the energy consumption of the DCPI [31]. Figure 21 shows the modeled response of a DC PUE curve as a function of the outdoor air temperature with and without an economizer mode. However, both the outdoor temperatures during which the economizer mode can be used and the resulting drop in energy consumption by the DCPI are highly dependent on the overall DC design and location.




Figure 21. Model illustrating the effect of an economizer mode on the PUE of a DC with respect to the outdoor temperature [17].

3.2.1 Dry Side

The most basic form of dry-side air economization is the direct use of outside air to cool the data center. When the outside temperature is low enough, dry-side economizer modes completely bypass the mechanical chiller and no supplemental cooling is necessary. However, the direct use of outside air in the DC has significant drawbacks related to air filtration, humidity control, and air handler fans size requirements [85]. There are two main methods used to perform dry-side air economization without the direct use of outdoor air: standard heat exchangers and heat wheels [126], [127], [128], [129]. An air heat exchanger isolates the inside DC air from the outside air traveling through the cooling system, with the thermal energy being transferred through the separating medium of the heat exchanger. This can be done with a variety of physical configurations and the most common, shown in Figure 22, uses an array of pipes containing the DC exhaust air with the outside air flowing around them.²

² The diagram also shows the addition of supplemental evaporative cooling. This technique lowers the temperature of the outside air in a manner similar to the foggers discussed in Section 3.4 while also wetting the outside of the pipes with the excess water to draw more heat away from the DC exhaust air as that water evaporates, the benefits of which are discussed in Section 3.2.2.




Figure 22. Indirect air-side economizer with supplemental evaporative cooling [130].

A heat wheel is a novel form of an air heat exchanger that a couple of companies have used successfully (Figure 23). A heat wheel is typically filled with a porous, corrugated aluminum medium and rotates very slowly, at approximately 4 – 12 RPM, requiring relatively minimal maintenance [131]. The inside and outside air pass through different sections of the wheel at any given time with little-to-no mixing. The rotation of the wheel through the DC exhaust air and into the incoming outside air transfers the heat to the outside air and out of the DC, with a heat transfer efficiency of approximately 72%.




Figure 23. A simplified schematic of how a heat wheel can be used to provide free cooling to a DC [126].

One of the main vendors of heat wheel cooling systems promises up to 85% energy savings with its modular systems compared to other cooling methods [132]. Some of the larger DC companies have had success themselves: Dell has achieved a PUE of 1.2 at its Tier 3, Class A1, enterprise-size DC in Quincy, Washington using a heat wheel [131]. Over the course of the year the DC uses only the heat wheel 75% of the time, a combination of the heat wheel and a mechanical chiller 17% of the time, and only the chiller for the remainder of the year.

3.2.2 Wet Side

Wet-side economizer modes are a bit more complex, but because of the higher heat transfer capacity relative to air (3,500 times higher [133]) there are significant opportunities to realize energy efficiency gains. The three main sources of Free Cooling for wet-side economizer modes are using either a large external body of water, direct outside air, or evaporative cooling with a cooling tower [126]. The external water or air is used to pre-cool the water used by the CRAH to cool the DC air, either supplementing or completely replacing the chiller.




Figure 24. Wet-side economizer mode based on evaporative cooling [126].

Evaporative cooling, shown in Figure 24, utilizes the latent heat required for evaporation to transfer the DC heat to the outside air using a cooling tower. This method is particularly effective as it requires 100 calories to warm one gram of water from freezing to boiling (0-100 °C), but it requires 539 calories to actually vaporize the water [133]. The cooling tower is also usually already present as part of the chiller system. As the water is sprayed into the tower, the molecules that vaporize draw a significant amount of energy from the surrounding molecules, lowering the temperature of the remaining water at the cost of having to continuously replenish the supply. Using a liquid heat exchanger, typically a plate-and-frame configuration, the lower temperature water returning from the cooling tower pre-cools the CRAH water and potentially allows the chiller to be bypassed. Google makes use of this configuration for some of its DCs [134]. Intel reported an 84% improvement in their HVAC effectiveness (IT power over HVAC power, only) using this type of system [85]. Another configuration uses the cooling tower return water to pre-cool the DC exhaust air before it enters the CRAH, either allowing the chiller to be turned off completely or to run at a lower load. Intel expects a 104% increase in the HVAC effectiveness for the system. This last example does not strictly fit under the definition of a wet-side economizer mode as presented here, but is referred to as both wet-side and dry-side by different groups.

3.3 Power Path Optimization

There is a significant difference between the type of power available on an electric utility grid and that used by most electronic components. Most grids provide 480 or 600 Vac electricity [135]; however, computer processor core voltages are on the order of 1 Vdc [136]. The transition between these two extremes requires a significant number of conversion steps, each one not 100% efficient. Currently, the majority of DCs convert the incoming 600/480 Vac to the 208/120 Vac accepted by the IT equipment. However, some modern servers are able to accept up to 240 Vac. Driven by the need to deliver higher power levels to the servers, using higher input voltage servers can result in DCPI energy efficiency gains by removing or decreasing the number of power conversion steps undertaken by the UPS, PDUs, and switchgear. These gains could be erased, however, if the power conversion steps in the server from the higher, 240 Vac power down to the eventual 1 Vdc power are not as efficient as the original conversion from 208 or 120 Vac down to 1 Vdc.

Although the entire conversion from 600/480 Vac all the way down to 1 Vdc is unavoidable, some DC power distribution networks have what could be additional, unnecessary conversion steps, as well. Traditional UPS systems typically have some sort of Vdc energy storage built in that needs to be able to charge and discharge. As a result, the 600/480 Vac is often converted to

Vdc to charge the energy storage and then converted back to 600/480 Vac before heading to the PDUs for conversion down to 208/120 Vac (Figure 25) [135]. A large number of conversion steps has three potential effects on a DC: increased electrical losses and resulting cooling load contributions, increased electrical utility power requirements, and increased floorspace requirements. Industry has developed multiple strategies to reduce the number of discrete conversion steps and increase the overall efficiency of the electrical power conversion, a couple of which show significant promise.




Figure 25. A traditional 480 Vac to 208 Vac power distribution chain for a DC [137].

Some companies have developed forms of local energy storage using batteries to reduce the number of power conversion steps in a relatively inexpensive and painless manner. One method is an integrated power supply and battery combination for use with blade chassis systems [138]. This allows for the elimination of the DC UPS, whose purpose is to maintain a consistent and high quality flow of power to the servers, because now the integrated battery has taken over this role. The integrated power supply and power combination can result in up to a 15% improvement in DC PUE and a 25% reduction in footprint. Another method puts back-up batteries directly on the server racks as well, eliminating two AC/DC conversion stages [139]. Combined with the use of high efficiency power supplies and voltage regulator modules, the upgrade has an estimated 25% annual energy savings per server.

More drastic measures have also been tested, with experimental systems showing significant potential for energy efficiency gains. Lawrence Berkeley National Labs demonstrated two novel DC power distribution systems with Vdc power: a facility-level configuration with the 480 Vac converted to 380 Vdc before the power enters the facility and a rack-level configuration with the 208/120 Vac converted to 380 Vdc. Both systems used Vdc powered servers instead of the standard 208/120 Vac powered servers [135]. The measured results showed a 7% and 5% decrease in the required energy, respectively, when compared to state-of-the-art power distribution systems with a 79% power system efficiency. At the time, the average Industry power system efficiency was 61%, resulting in an estimated 28% decrease in energy requirements.³

Another configuration is a row-based UPS which accepts a 400 or 480 Vac input and outputs -48 Vdc to the racks [140]. The implementation of such a system is shown in Figure 26. A Green Grid study showed that this type of system could be around 90% efficient [141]. When compared to a legacy 480 to 208 Vac system it was approximately 25% more efficient. However, a contemporary 480 to 208 Vac system was also approximately 25% more efficient than the legacy system. Although not evaluated in the study, the smaller power system footprint is likely the

³ Note that these results did not include the cost of cooling the DC and are not directly analogous to the DC PUE, but only one component.

most significant advantage of Vdc-centric power distribution systems. Estimates of this benefit range as high as 50% less floor space than comparable Vac systems [137].




Figure 26. A row-based Vdc UPS used to convert the 480 Vac utility grid electric power to the -48 Vdc server input [137].

In general, the work mentioned thus far claims that reducing the number of discrete power conversion steps from the Vac utility power to the Vdc power required by the electronic components will result in efficiency gains. However, an independent study showed that simply upgrading the components of a legacy DC power conversion and distribution chain can eliminate most of those gains; i.e., a modern traditional power conversion and distribution system can be just as energy efficient as a modern Vdc-based system. As a result, the consensus benefit of using any of the aforementioned technologies (local energy storage, facility-based Vdc power distribution, row-based Vdc power distribution, etc.) is that a significant amount of floorspace can be freed up with some of these methods, which is not insignificant for many applications and might be worth the capital required for the conversion.

3.4 Evaporation Cooling using Foggers

Humidity is a significant concern for a DC operator. In addition to specifying recommended temperature ranges, ASHRAE TC 9.9 recommends humidity ranges, as well [16]. The 2011 recommendations say that a 30-60% relative humidity range should be maintained in a DC. Going outside of this range risks accidental electrostatic discharge or excessive wetting of the equipment, which can lead to corrosion [142]. Relative humidity is the percentage of water vapor in the air compared to the maximum amount it can hold, which increases exponentially as the air temperature increases. A lot of effort goes into increasing the relative humidity of the air before it goes back into the DC as cooling units have the effect of removing moisture from the air due to condensation. Keeping in mind the cooling efficiency benefits of evaporative cooling, discussed in Section 3.2.2, some evaporative cooling techniques can simultaneously provide cooling and help maintain the desired humidity levels, thereby further improving the overall DC efficiency by eliminating the need for supplemental humidifiers. The focus of this section is evaporation cooling using foggers, although wet meshes can also be used in a process that is analogous to the swamp coolers used for residential cooling in arid climates.



Figure 27. An evaporation cooling fogging system [133].

Evaporation cooling using foggers sprays fine mists of water into the air being sent into the DC [126], [133]. The water travels along pipes and is forced through nozzles, separating the molecules as they enter a significantly lower pressure and higher volume environment. The change in pressure causes the water to vaporize. It is key to understand that this is an adiabatic process and no additional heat is added to the water or the air to cause the evaporation. The water draws energy from the surrounding air molecules to make the transition, lowering the temperature of the air while also increasing its relative humidity. When used with an economizer mode, this can increase the temperature range during which the traditional chiller does not need be turned on. A case study for a DC in Oregon, which used a dry-side economizer with direct air, showed that their fogging system could increase the economization periods and was predicted to reach a PUE of 1.15 [133]. Foggers can be used in systems without economizer modes, as well, lowering the high temperature return air before it enters the cooling unit. In general, a high-pressure fogging system with a 10 HP motor is predicted to be able to replace a 500 ton traditional chiller. This, of course, is highly dependent on the local climate [143].

3.5 Close-Coupled Cooling

Close-coupled cooling comes in four main forms: in-row air conditioners, rear door heat exchangers, overhead air conditioners, and in-rack closed-loop cooling [144]. In-row air conditioners take the traditional large-scale cooling units and separate them into smaller units. The main benefits of this approach are modularity and proximity: modularity allows for additional cooling capacity to be easily installed along with new IT capabilities and proximity helps to increase cooling efficiency. Reducing the distance the exhaust air has to travel reduces the necessary fan size and, therefore, electrical power consumption. Also, as each row has its own fan, the speeds can be adjusted to more closely match the local cooling loads, allowing their power consumption to be more effectively minimized. A combined CFD study and economic analysis of a 1.2 MW DC using both water- and air-cooled chillers showed that while in-row air conditioners would have a greater first cost, the lower operating costs cause them to be more cost effective after only 3 years [145]. Overhead air conditioners are designed to maximize available floor space by placing the air conditioning on top of the rack. Instead of the air traveling up from the rack and into an exhaust plenum, most or all of the heat has already been removed before the exhaust air re-enters the room. It has very similar benefits to the in-row air conditioners, but can be expensive to install on a large scale and have a long payback period [145].

Both overhead and in-row air conditioning systems are often used to mitigate hot spots in a traditionally cooled DC to supplement the room-level cooling unit. When used in lieu of larger systems, the major trade-off to note is that while the operator gains modularity and flexibility by using more, smaller fans, larger fans have greater maximum efficiencies [146]. The operator might be able to achieve greater system-level efficiency by selectively turning on and off different fans as they are needed, but the individual fans themselves will have a lower maximum efficiency. Also, decreasing the distance between the cooling unit and the servers has a positive effect on the inlet coolant temperature [144]. Decreasing the distance the air travels raises the maximum acceptable air temperature for the cooling unit and by extension the maximum acceptable cooling unit inlet water temperature. This greatly increases the opportunities for economizer modes while decreasing the amount of work the cooling unit must perform.

Rear door heat exchangers move the cooling coils from inside the centralized cooling unit to the rack itself [144]. As most IT equipment is designed to move air from the front to the back, placing cooling coils directly behind the IT equipment means that some or all of the heat can be removed as the air leaves the rack. This method takes up very little floor space in the room as the coils are mounted in the rack itself. Some companies also sell modular systems, allowing the capital cost to be spread out as the DC IT capacity is expanded [147]. In addition, some rear door heat exchanger implementations, referred to as passive doors, do not require supplemental fans beyond those already on the server or blade chassis. However, passive doors can require each rack to be better sealed against unintended airflow while also simply keeping the localized airflow work on the servers themselves, not necessarily reducing it overall. Rear door heat exchangers can also bring high performance refrigerant directly to the racks, which increases cooling efficiencies because of their greater ability to transport heat relative to air or water. Pipes, pumps, and controls are still necessary and must be installed somewhere, and the equipment and installation tradeoff becomes increased piping and coolant pump requirements against decreased air movement requirements. In-rack closed-loop cooling systems combine all of the IT, power distribution, and cooling equipment in what are essentially self-contained DCs on the rack scale. This type of system benefits from being extremely modular with very small individual building blocks and is typically used in unique applications requiring small levels of IT capability.

3.6 Summary and Recommendations

Although any upgrade to a DC to increase its energy efficiency will carry a substantial up-front capital cost, a couple of technologies were found that promise to be significantly less expensive than those that will be discussed in Section 4. The first, that should always be considered, focuses on maximizing the efficiency of air cooling in the DC room with basic airflow management technologies, often referred to as virtual containment. Sealing unnecessary openings in a raised floor and in the racks, for example, often allows the cooling unit set points to be raised, reducing costs with more energy efficient cooling. If virtual containment is insufficient, hot/cold air containment technologies can be investigated. There is a wide range of products available to perform hot/cold air containment in the DC, which can improve the air cooling effectiveness by as much as 166%. Some low-cost and easy-to-install methods perform just as well as the higher-cost methods, providing DC operators with a range of effective options based on their specific needs. Most importantly, after any new airflow management technology is implemented the ASHRAE TC 9.9 temperature and humidity guidelines and individual

equipment manufacturer specifications should be consulted to allow the DC operator to decrease the cooling system output as low as possible to maximize the energy efficiency gains.

In addition to increasing the efficiency of the cooling system, upgrading and updating the power distribution system often produces significant efficiency gains. Although there are some companies espousing a move toward a more Vdc-based power distribution system, independent experiments have shown that simply modernizing a legacy system with more energy efficient equipment can increase the overall power distribution system efficiency by as much as 25%. While this would be by no means inexpensive, it also has the attraction of being tied to a facility upgrade as opposed to a DC upgrade, placing the responsibility in the hands of the organization that would stand to benefit the most as it typically pays the DC electricity bill.

The next three technologies represent a larger change and likely a larger capital investment to implement: supplemental evaporative cooling using foggers, cooling system economizer modes, and close-coupled cooling. For the cost of a system of high pressure water pipes, nozzles, and pumps, it is possible to take advantage of the significantly better heat transfer capabilities of evaporative cooling with DCs that use air at any point in their cooling system. The water can be evaporated into outside air used for Free Cooling or into air used in a closed-loop flow to-and-from the DC room, for example. The impact of these systems depends greatly on the climate where the DC is located. The same applies to the use of Free Cooling; however, economizer modes aim to improve its potential in less-than-ideal environments. By using low-temperature air or water from the environment to supplement the traditional cooling system, Free Cooling can be cost effective in a larger number of cases.

Close-coupled cooling has started to gain in popularity due to its ability to accommodate the higher IT power densities of blade servers. This has driven down the price of the systems and made it easier to take advantage of the cooling efficiencies granted by closer proximity to the original source of the cooling load as well as their greater modularity. Row-based close-coupled cooling is currently used by at least one major company, but modular rear-door heat exchangers with high performance refrigerant represent for many the next logical step in DC cooling systems past room-based cooling units because of their ability to bring the cooling system directly to the rack. Overhead systems are more expensive than in-row cooling and rear-door heat exchangers, but are often chosen in DCs where floorspace is at a premium. All of these methods would simplify DC design and operation by decreasing the importance of the DC floorplan and CFD airflow studies. While likely more expensive to implement than some of the other technologies discussed in this section, particularly airflow management, these close-coupled cooling systems are still expected to be significantly less expensive than those discussed in Section 4.

A couple of these technologies also hold promise for DoD unique applications. Three separate examples were given in Section 1.1 for unique DoD DC applications: the CPP RWS on a HMMWV, the E-8C JSTARS, and aircraft carriers. For all of these systems, major C4 capabilities (the DCs) were added onto a separately designed platform as a payload, can be expected to expand significantly as tactical computing needs increase, or both. This often results in a very inefficient DC as the goal during its installation or expansion is to get it to work and rarely to have the whole resulting system (including the platform) be energy efficient. Two steps are highly recommended to help achieve an energy efficient C4 platform: implementing hot/cold air isolation for the cooling systems and investigating opportunities to integrate Free Cooling. Free Cooling could potentially be very advantageous for sea-based platforms and high altitude airborne platforms as they both are surrounded by very abundant sources of low-temperature water or air, respectively.

4 High-Capital Technologies

All of the technologies mentioned thus far are expected to not require a significant capital expense to either upgrade an existing DC or integrate into a new facility. The ones discussed in this section, however, can require significant infrastructure and are typically considered only for new facilities or large-scale upgrades. The first set of technologies, direct-to-chip cooling, represent a new method of improving the efficiency of DC cooling by removing the heat as close as possible to the servers themselves. Next, many companies have tried to increase energy efficiency by pushing the system boundaries beyond just the DC facility and making use of the DC waste heat for other applications. While this will not decrease the DC PUE, it can reduce the energy requirements of other nearby systems. Lastly, there are efforts to harness waste heat from other co-located systems and apply it to the DC facility. This typically comes in the form of an on-site Co-Generation (CoGen) or Tri-Generation (TriGen) plant.

4.1 Direct-to-Chip Cooling

Direct-to-chip cooling has the potential to remove produced heat in a very energy efficient manner and also enable higher computing densities. This is due to the minimization of the distance between the coolant fluid and the heat source, as well as, the ability to use liquid and two-phase cooling techniques (two-phase direct-to-chip cooling is discussed in Section 5.2). The major hurdles preventing these methods from being commercially viable are the large capital cost, which is higher for both the DCPI and the IT equipment, and the significantly greater chance for leaks. There are a couple of examples of research facilities constructed to investigate the maximum achievable DC energy efficiencies: IBM's Aquasar prototype at ETH Zurich [148] and the NW-ICE supercomputer at the Pacific Northwest National Laboratory Energy Smart Data Center testbed facility [149].

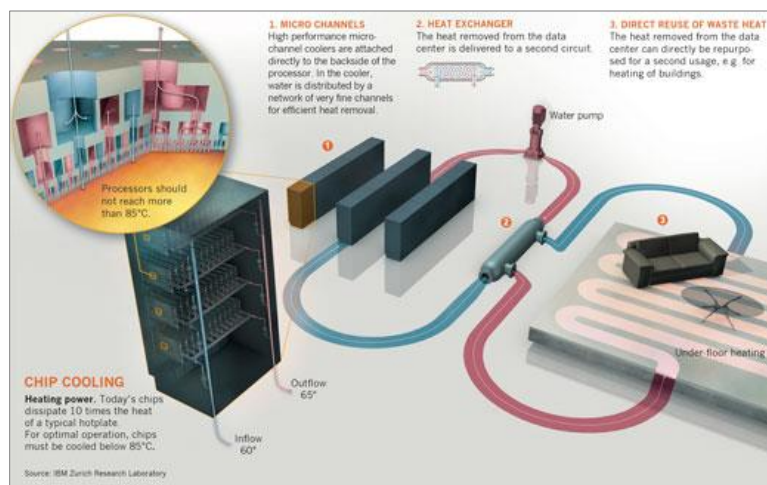


Figure 28. A direct-to-chip cooling prototype system with waste heat reuse [148].

The IBM Aquasar system (Figure 28) uses microchannel coolers mounted directly onto the chips in the servers. The inlet temperature for the coolant fluid can be as high as 60 °C with the outlet fluid temperature averaging 65 °C. IBM claims that this system consumes up to 40% less energy

than a comparable air-cooled system. Also, because of the high temperature waste heat returning from the IT equipment, a heat exchanger is used to transfer the heat to an underfloor liquid heating loop in adjacent buildings on the university campus. Figure 29 shows an example of the special blade servers required to implement direct-to-chip cooling. This system is able to directly collect up to 85% of the heat produced by the blade [150].



Figure 29. Example of a modified blade server with the hardware required for direct-to-chip cooling [148].

4.2 Waste Heat Reuse for Space Heating

In a traditional DC the energy used by the IT equipment is essentially lost, used by the electronics and converted into the cooling load that is rejected to the outside environment by the cooling system. A number of companies have looked to leverage that waste heat because, even with the most energy efficient DC, a 1 MW IT load will still produce at least 1 MW of waste heat with current computing technology. The uses for the waste heat vary, but the viable applications for each site depends heavily on the waste heat quality (or temperature) and therefore the cooling system type. Table 6 provides an overview of the typical and maximum temperature waste heat available using air, water, and two-phase cooling methods. As mentioned previously, blade servers can provide higher return coolant fluid temperatures and, therefore, can be expected to produce waste heat closer to the maximum end of the range, whereas “xU” servers can be expected to produce waste heat closer to the “typical” values shown in the table. In the next two sections, both waste heat reuse for space heating and water desalination are discussed. Space heating is discussed because it is the most common use with DCs. Water desalination using waste heat from a DC is discussed in Section 5.1.

Table 6. The typical waste heat quality available using three main DC cooling techniques: Air, Water, and Two-Phase [151].

Cooling Method	Typical Waste Heat Quality (°C)	Maximum Quality (°C)
Air	30	45
Water	60	75
Two-Phase	75	90

Space heating is a very common use for DC waste heat because it has very low temperature requirements and can even be performed with the 30 – 45 °C heat typical of air cooling, the most prominent cooling method used to cool a DC. IBM built a DC in Switzerland where the waste heat can be used to heat water for a nearby town's pool, which was claimed to be enough heat for 80 homes [152]. In this arrangement, the townspeople are able to make use of the waste heat at no cost but contributed a portion of the initial capital costs. On the other end, IBM recoups some of the initial capital expense of building or upgrading the facility, sometimes a big factor preventing the use of more energy efficient technologies.

Multiple companies have also entered into agreements with local towns or utilities to sell the waste heat for district heating. Telehouse West announced a 130,000 sq. ft. DC in the Docklands of London that could provide up to 9 MW of heating power to the district heat network [153]. Academica announced a DC under a cathedral in Helsinki, Finland that would input waste heat into the local district heating to supply heat for up to 500 large private homes. The company said at the time (2009) that the deal would be expected to trim \$561,000 from the annual power bill due to the deal [154]. As these examples illustrate, even with older, air cooled DCs, district heating with DC waste can prove cost effective.

More advanced and efficient cooling technologies can increase the gains of waste heat reuse for space heating. The IBM Aquasar system mentioned in Section 4.1 provides approximately 6 kW of waste heat to the nearby campus buildings from a 6 Teraflop, 13.3 kW rack cooled using direct-to-chip cooling [148]. Scaling up the technology to an entire DC with dozens or even hundreds of racks would provide a significant amount of waste heat. Waste heat reuse is also compatible with free cooling: the TelecityGroup DC in Paris mentioned in Section 3.2 uses its waste heat for an on-site 'Climate Change Arboretum' [123].

4.3 Co-Generation and Tri-Generation

Both co-generation (CoGen, sometimes referred to as Combined Heat and Power, CHP), and tri-generation (TriGen) have two important characteristics in common: on-site power generation and the use of absorption chillers. CoGen uses the waste heat from the on-site power generation to power an absorption chiller, providing the DC with both electric power and cooling. TriGen is similar to CoGen except that it also makes use of surplus waste heat for space heating.

Absorption chillers are relatively complex, but the basic premise is that the working fluid is evaporated at a high pressure using the heat source, condensed and transported to a low pressure compartment, often an external cooling tower, where it is evaporated over the chilled water return [155]. Figure 30 shows a simplified schematic of a single-effect absorption cycle.

Lowering the pressure of the working fluid (dark blue) decreases the vaporization temperature, allowing it to draw its latent heat of vaporization from the relatively low temperature chilled water return (light blue), lowering the chilled water temperature even further before re-entry into the CRAH.




Figure 30. A simplified schematic of a single-effect absorption cycle [156].

CoGen has been a focus of the U.S. Government Environmental Protection Agency (EPA) as a method to reduce energy needs of DCs [157]. An example facility installed by Qualcomm at its headquarters in 1995 has both a 2.4 MW and 4.5 MW (2005 facility upgrade) natural gas powered turbine on-site. The waste heat from the turbines is captured and either used to produce hot water for absorption chillers or used to directly power a 1,400 ton absorption chiller, in the case of the newer 4.5 MW turbine. The site had an estimated \$775,000 in yearly savings stemming from a reduction in annual operating costs (\$500,000), hot water production (\$100,000), and power utility demand (\$122,000). A more recent CoGen facility was installed by KPMG at its northern New Jersey DC [158]. The facility requires 22% less fuel than a more traditional grid-tied power system and has an expected 5-7 year payback period for an original \$2.4 million upfront capital cost, taking into account a \$900,000 state tax credit given to the company.

IBM and Syracuse University teamed up to implement a TriGen system on the university campus, shown in Figure 31 [159], [160]. It uses on-site natural gas microturbines, whose 307 °C turbine exhaust is used both for space heating and to power two absorption chillers. The chilled water exits the absorption chillers at 7 °C and is used in rear door heat exchangers to remove the server waste heat from the racks.

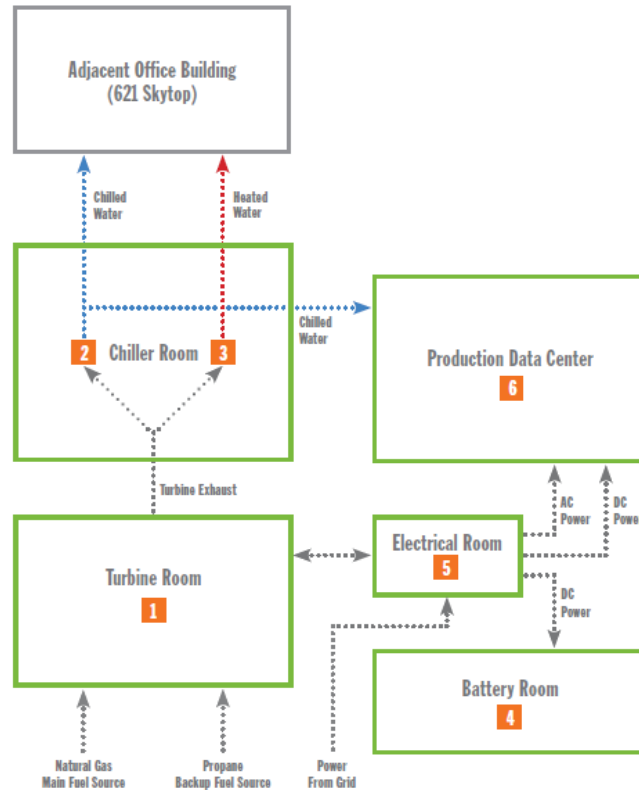


Figure 31. A simplified operational schematic of the IBM Green Data Center at Syracuse University [161].

CoGen and TriGen both have the potential to enable an energy efficient and energy secure DC, although at a significant capital cost. In terms of increasing DC energy efficiency, CoGen and TriGen can reduce the transmission losses associated with getting power from the power plant to the DC. Since the electricity does not need to travel across the grid, it can be generated at voltage levels significantly closer to those used by the DC, resulting in fewer conversion losses. For example, if the power plant directly outputs 208 Vac, then the power does not need any additional conversion steps before it can be sent to a traditional server. Also, the methods make use of the waste heat from the now on-site power plant to provide DC cooling. Absorption chillers powered by power plant turbine exhaust use free energy in the form of heat to drive their cooling process.

A similar setup to CoGen and TriGen using fuel cells has been implemented at one large DC, taking advantage of waste heat from off-the-grid power generation for greater energy efficiency and resiliency [162]. A 5 MW waste heat recovery plant is attached to the natural gas compressor station that fuels the on-site fuel cells supplying electricity to the DC. The DC was predicted to achieve a PUE of 1.2. In general, fuel cells have been explored by many large DC operators because of their ability to provide more efficient, on-site, power generation compared to grid operation. However, they are just starting to be used as the sole DC electric power at some facilities and previously have been used for backup power generation as an alternative to the more traditional diesel generators [163].

4.4 Summary and Recommendations

The technologies described in this section can be thought of as potentially representing the next generation of advanced commercial DCs. Direct-to-chip cooling promises to significantly increase the heat transfer efficiency and, by extension, the DC computing density. CoGen and TriGen systems are attractive because the DC utility requirements often are too large to be reliably supplied by the commercial grid. As single DC facilities continue to grow in size and their missions require a greater level of IT availability, investing in CoGen and TriGen will become more and more attractive and cost-effective. Also, as DC IT densities and heat transfer efficiencies increase, the opportunities for using the waste heat produced by the servers and other pieces of equipment will increase. While current applications are limited to local district heating, desalination, water filtration, and absorption cooling systems are very attractive future uses.

As these technologies advance further, the DoD should monitor their progression and investigate their potential when significant installations or upgrades are scheduled for new or existing facilities. A couple of technologies warrant more immediate attention. Consolidating DoD DCs means that a higher degree of resiliency is necessary as more mission critical operations are centralized in a single system. Implementing a CoGen system for a large-scale, mission critical DCs will allow them to be isolated from vulnerable electric utility grids. Large-scale platforms carrying significant DCs can implement their own, mobile, CoGen systems as well. If the waste heat from the platform prime power generation system is used to drive an absorption chiller that cools the onboard DC, the platform's overall energy efficiency increases from reduced auxiliary electricity demand. Aircraft carriers, for example, already use the waste heat to power onboard water desalination systems. Besides investigating whether there is sufficient waste heat to power an absorption chiller as well, the Navy should investigate whether the DC waste heat can be used to supplement the prime power waste heat to potentially increase the desalination capacity.

5 Nascent Technologies

During this review, multiple technologies were found that are not mature enough for the DoD to currently leverage for their DC needs. These technologies, whether being investigated by Academia or by Industry themselves, are summarized in the following sections.

5.1 Filtration and Desalination

For some DoD DC applications there is a large supply of sea water close-at-hand, and cost effective desalination is very desirable and challenging. Academic studies have been performed showing the feasibility of using multiple effect distillation (MED) systems with a minimum of 75 °C waste heat [151]. In one specific study by Li *et al.*, 90 °C inlet water is used with a model of a hybrid absorption heat pump system to desalinate up to 4,400 gallons/day of sea water through a model of a four-stage MED system [164]. For this system, the 90 °C inlet water left the last stage of the desalination system at 27 °C, potentially eliminating the need for supplemental cooling before re-entry into the DC cooling system. Although not a commercialized technology, there is at least one company working in the area of water desalination that warrants investigating: Trevi Systems. They use the forward osmosis (FO) process (Figure 32) driven with 75 °C waste heat for ocean desalination, requiring approximately 50 MJ/m³ while operating at ambient pressures [165].

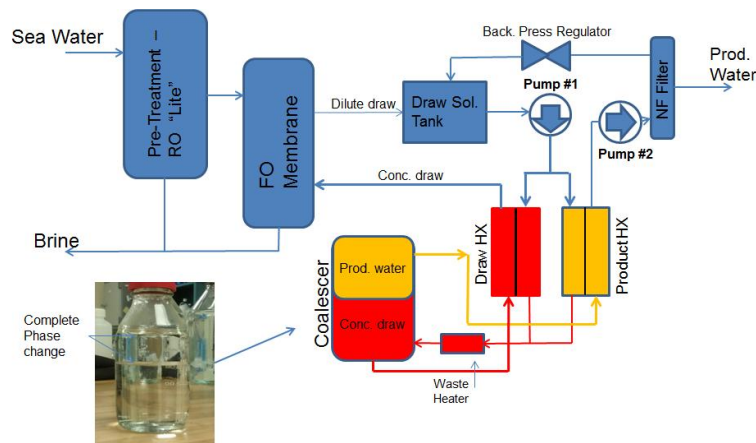


Figure 32. A schematic of the Trevi Systems forward osmosis (FO) desalination process [166].

5.2 Two-Phase Direct-to-Chip Cooling

At the Pacific Northwest National Laboratory Energy Smart Data Center their NW-ICE compute rack uses 12 kW and is able to operate using SprayCool™ modules with Fluorinert™, a range of coolant fluids produced by 3M™, as the coolant [149], [167]. The modules are fully enclosed and mist Fluorinert™ onto the surface of the processor, using evaporative cooling to remove the heat from the surface. The heat in the Fluorinert™ vapor is then transferred through rack-based heat exchangers to the facility-level chilled water. Also, because of the use of Fluorinert™ as a two-phase coolant (available with boiling points ranging from 30 – 215 °C) the extracted heat can be at relatively high temperatures for waste heat reuse. As with the previous example, the SprayCool™ system requires specialized servers with channels for the Fluorinert™ in addition to the modules that mount on top of the processors, replacing the traditional heat sinks.

5.3 Immersion Cooling

A handful of companies are developing and marketing DC cooling solutions that completely submerge the IT hardware and power distribution systems in an electrically insulating liquid, aiming to take full advantage of the up to 3,500 times greater thermal capacity of these liquids compared to air [133]. Unlike the rear-door heat exchangers discussed in Section 3.5, the liquid is brought into direct contact with the server's electronic components. These techniques differ from the direct-to-chip cooling techniques discussed in Sections 4.1 & 5.2 because the entire server is in contact with the liquid, as opposed to only specific components connected to the network of coolant-carrying channels in the server. There are two main categories: single-phase and two-phase immersion cooling. Single-phase immersion cooling relies on pumps to circulate the cooling liquid out of the environment with the IT equipment where the heat is removed with a liquid heat exchanger [168], [169], [170]. Either the entire rack of IT equipment is laid down on its side in a liquid bath or the servers are designed to contain their own liquid bath. Using environmentally friendly, non-flammable, and odorless fluids, the vendors advertise an up to 95% decrease in cooling energy requirements. There are a range of configurations available and some are offering modular, containerized systems as small as 8 blade servers and as large as an ISO shipping container. Two-phase immersion cooling relies on the evaporation of the cooling liquid when it comes in contact with the IT equipment, removing the heat as it floats to the top of

the immersion bath to condense on a secondary coolant coil to remove the heat completely from the environment [171]. The fluid has a useful lifespan of 30 years, and the two-phase immersion cooling system only requires a power and internet hookup, offered in individual racks/tanks or as 6 tanks in an ISO shipping container.

5.4 Autonomous Data Center Management

A DC represents a very complex environment as each different server can have a different load, and the flow of air around them can shift drastically based solely on how the load is distributed. To address this, many companies market power and temperature monitoring hardware and software with the aim of enabling close to real-time DC operations optimization by a human DC operator [39], [172]. The area is starting to mature as the Energy Efficient High Performance Computing (HPC) Working Group [173], an Industry, Academia, and Government partnership, recently presented a proposed methodology for power measurements of large-scale high-performance computing systems at a conference [174]. The goal of instrumenting a DC and taking electric power, thermal, and other measurements is to use that data to inform the automatic adjustment of the DC systems to optimize operations for energy efficiency [70]. This, however, does not appear to integrate the control of the IT load distribution with the control of the DCPI. Recall that there are software tools capable of automating the powering-on and -off of specific servers based on the current IT load [93]. However, as the IT load distribution and the DCPI performance are inextricably linked, any complete autonomous DC management system must be able to incorporate and control both. There are innovative companies working in this space, integrating automatic software management with automatic infrastructure management [175], [176]. In one case study the company reduced energy costs 30% by automating the nightly shutdown of the DC when it is not being used [177]. Implementing a reservation system that automatically powers up the necessary portions of the DC before scheduled use is predicted to increase the energy cost savings to 75%. The goal of this technology is to adjust the applications, IT hardware, and DCPI all in real-time according to existing or predicted IT loads to optimize DC energy efficiency. The development of the underlying technology warrants close monitoring as, once it is mature, there will likely be significant opportunities for DC energy efficiency optimization and increased resiliency due to the ability to adapt to failures, variability, and uncertainty.

5.5 Integrated Data Center Thermal and System Modeling

As DCs become more complex and their missions more critical, it is advantageous to model their performance before facilities are built, systems are upgraded, or new applications are installed. The performance of a DC is directly related to two very different and complex subsystems: the thermal transport behavior and the computational performance of the applications, etc. There are multiple companies offering CFD modeling of the thermal behavior of the DC, typically for air-cooled DCs [59], [60], [61], [62]. This allows for the planning of the DC layout and even the distribution of equipment and maximum workloads within a rack. The modeling of the computing performance of a DC is not as mature, however. Intel markets its CoFluent Studio software [178], and some researchers from Academic institutions, in collaboration with a corporate partner - Tata Research Development and Design Center, developed their Modellus system [179]. However, to actually predict the behavior and performance of a DC to loads or environmental conditions requires the interaction of the two different types of models.

Completely integrated DC models would also be invaluable for informing the autonomous DC management algorithms discussed in Section 5.4.

5.6 Thermoelastic, Magnetocaloric, and Electrocaloric Cooling

Mechanical refrigeration is based on the vapor compression cycle, and, in addition to being inefficient, there is often the potential for the accidental release of harmful vapors to the environment. There are three other forms of heat transportation that are being investigated. They are still in the proof-of-concept phase, and all are aiming to be more efficient, have fewer moving parts, and be more environmentally friendly than current systems. Thermoelastic cooling has been a focus of the DOE and ARPA-e, making use of the latent heat transferred during the solid-to-solid martensitic phase transformation of thermoelastic materials, such as shape memory alloys (SMAs) [180], [181]. Thermoelastic cooling is projected to have approximately double the coefficient of performance of a vapor-compression refrigerator. Magnetocaloric cooling systems operate based on a centuries-old premise; however, recent breakthroughs in permanent magnet and magnetocaloric materials are spurring the development of new prototype systems that operate at room-temperature [182], [183], [184]. Magnetocaloric cooling operates around the increase and associated decrease in temperature of a metal as a magnetic field is introduced and removed, an adiabatic process. To achieve strong enough magnetic fields, superconducting magnets had been required until recently and removing the requirement for cryogenic cooling has significantly increased its feasibility. Electrocaloric cooling operates in a manner very similar to magnetocaloric cooling; however, in this case, an electric field is applied to the material as opposed to a magnetic field [185], [186]. Although the theoretical predictions show that efficiencies could range from 60-70%, there are significant challenges still to be overcome as the greatest efficiency performance is seen from thin films and therefore not practical for high thermal mass transfer applications.

5.7 Piezoelectric Fans

In addition to a trend toward more power dense computing electronics there is a concurrent trend toward smaller systems. As with increasing the power density, decreasing the overall size of computing systems is limited by the ability to remove the heat. Traditional axial fans do not scale down well, particularly to sizes in the tens of millimeters, and microchannel cooling is expensive. It is within this space that researchers have investigated the use of piezoelectric materials to create forced convection currents [187], [188]. The most basic configuration is the application of the piezoelectric material at the base of a cantilevered beam, which can then be oscillated back-and-forth with the application of a Vac signal. Although they have the benefits of low cost, low power consumption, and low noise, there are some significant challenges that need to be addressed. Piezoelectric cantilevered beams require very high actuation voltages and have high actuation frequencies. As a result, their actuation is nontrivial, particularly on a small scale. Also, some of the most effective piezoelectric materials are typically ceramics and therefore have low yield stresses and cycle limits. It is expected that significant work will continue toward improving the performance of piezoelectric materials as they are also very often used for precision linear actuation and vibration energy harvesting.

5.8 Jet Impingement

One method being explored to potentially produce a higher rate of heat transfer than two-phase direct-to-chip cooling is jet impingement. Jet impingement is the spraying of a small diameter jet of fluid onto the surface of the material or chip, resulting in very high heat transfer coefficients because of the very small boundary layer between the fluid and the surface [189] [190], [191]. Advanced Thermal Solutions has created a prototype cooling system, Therm-Jett™, for a 1U rack server using air as the jet fluid and was able to achieve a 20 °C improvement over conventional methods [192]. The primary challenge is developing a microchannel array of nozzles capable of cooling a large surface area relative to the ideal diameter of the jet for the application, which can be smaller than 1 mm [193]. The jet fluid needs to be distributed among the many different nozzles equally and the fluid needs to have a viable return path that will not interfere with any of the other jets. If these challenges are able to be overcome, the technology shows great promise for highly effective cooling on a small scale.

5.9 3D Stacked Chip Architectures

One method multiple researchers are investigating to further increase the density of computing systems is to stack processors on top of one another in a 3D architecture. Two European consortiums, are investigating two of the major challenges associated with the technology. FlexTiles is focused on defining the interconnections between multiple layers and developing an energy-efficient, programmable manycore platform prototype [194]. The platform will have a reconfigurable layer stacked on top of a traditional manycore layer. By altering how the reconfigurable layer uses the manycore layer, the chip can be optimized for each individual application. By being able to quickly and independently adapt to different demands, the chip will be capable of improving load balancing, decreasing power consumption, and improving resiliency to faulty modules compared to more traditional architectures. The CMOSAIIC project is focused on developing solutions for cooling multilayer chips [195]. It is meeting this challenge by investigating associated micro-fabrication techniques, the bonding of stacked layers, developing more advanced dynamic thermal modeling, and performing experimental testing. Although there are a couple of corporations involved in the research, it appears to still be in the R&D phase and bears watching as the technology progresses.

5.10 Heterogeneous System Architectures (HSA)

AMD has been devoting significant research toward developing a new programming platform that is able to operate the central processing unit (CPU) and graphics processing unit (GPU) in parallel [196], [197]. GPUs were traditionally used solely for specialized graphics-related tasks; however, they have become more and more generalized, and some applications are using them to perform parallel computing. As the pairing was not originally intended for parallel computing it is not straight forward or efficient to do so and is, therefore, limited to only the most demanding of applications. AMD is aiming to develop a new processor architecture and framework that allows for the easy parallelization of the CPU and GPU, opening up more processing power for more applications to take advantage of. This can significantly improve the power efficiency for DC operations as preliminary experiments have shown a 2.3x relative performance gain coupled with a 2.4x reduction in power requirements. To help create wide-spread adoption of the

technology the HSA Foundation was formed with ARM, Imagination Technologies, MediaTek, Texas Instruments, Samsung Electronics, and Qualcomm.

5.11 High Temperature Electronics

The drive for electronics able to operate correctly and efficiently at higher temperatures has been driven significantly by the oil and gas extraction sector to drill deeper and the automotive and aerospace industries to move more toward distributed electronic control [198]. For example, aircraft engine control electronics are necessarily bundled far enough away from the engine itself to ensure temperature maximums are not exceeded. Distributing the controls immediately on the engine reduces the mass, complexity, and failure points for the overall systems. Some of these demands have been met by commercially available Silicon-on-Insulator (SOI) technology, which can operate between -55 and 255 °C and greatly reduce leakage losses above 200 °C. The SOI devices currently available are limited to simple power switches, conversion devices, and 8-bit microcontrollers [199], [200]. Wide band gap materials, such as Silicon Carbide (SiC), are being investigated to take the next step toward greater operating temperatures and increased power efficiencies. Recent research has demonstrated simple SiC circuits at 500 °C lasting for thousands of hours, operating down to -125 °C, and having greater efficiency than SOI power electronics at temperatures >85 °C [201]. Significant challenges remain for pushing SiC devices past 100 discrete transistors. Developing interconnects and packaging able to withstand the greater temperatures are two examples. In terms of DC energy efficiency, the potential benefits are significant. High temperature electronics can achieve greater efficiencies while being able to tolerate higher DC temperatures. More efficient power electronics reduce the losses associated with stepping the electrical power down from the grid level (600 or 480 Vac) to that used by a processor (~1 Vdc). Higher operating temperatures reduce the necessary cooling load while also significantly increasing the range of environmental conditions Free Cooling and economizer modes can be used. Advances in the manufacture of high temperature SOI and SiC electronics should be watched as a possibility for improving DC energy efficiency.

5.12 Reversible Logic

In a DC it is traditionally thought that to operate as efficiently as possible the power consumed by the DCPI should be minimized, and that, in an ideal DC, it is zero. This logic, however, ignores that the biggest consumer of energy in a DC is the IT equipment itself. Rolf Landauer's principle is that erasing a bit of information always dissipates a small amount of energy, and with current logic gates this loss of energy is unavoidable for two reasons [202]. First, standard logic gates output "garbage states" that are never used. Second, as standard logic gates are not perfectly reversible, verifying a computation requires erasing bits and repeating that computation at least one more time. Reversible logic aims to allow computing which does not have to erase bits as the gates are both logically and physically reversible, thereby theoretically dissipating no energy as a result of Rolf Landauer's principle. This technology has very limited experimental success, however, and requires a significant amount of research to become practical for DCs [203].

6 Conclusions

DC energy efficiency has grown in importance as DC services continue to become more integral to daily life and military operations. There has been a lot of Independent Research & Development (IR&D) devoted to this challenge as companies strive to increase profits by reducing operating expenses. Based on this review of the current technologies and methodologies Industry is using to increase the energy efficiency of their DCs, there is significant potential for the DoD to leverage that technology with the aim of streamlining their own operations and improving mission effectiveness. To aid DoD organizations in overcoming the difficulty of funding the necessary capital investment for significant improvements, unique contracting vehicles, such as an ESPC, can be explored.

The first step for any DC should be a complete inventory of all IT equipment and applications, followed by the decommissioning of any found to be unused. Legacy applications stuck on outdated hardware should be transitioned to modern infrastructure as the simplest step is taking advantage of the improvements in IT hardware energy efficiency that have occurred recently. Modernizing the IT equipment is often overlooked as the Industry standard DC energy efficiency metric, PUE, does not incentivize decreasing the power consumption of the IT equipment itself. Putting equal weighting in procurement decisions to both the expected PUE and an Industry benchmark for server energy efficiency, such as the SPEC.org or the PAR^{4®} metrics, would significantly increase the energy efficiency of DoD DC operations. While more and more DCs are becoming virtualized, the ability to have virtualization software tools automatically power on and off individual servers does not appear to be widely known or used. Implementing the specific virtualization tools that dynamically adjust the active IT hardware based on the IT load will drastically reduce the amount of power consumed by unnecessary, idle servers.

Along these same lines, a traditional pitfall for every DC operator installing new capacity or an entirely new facility is the tendency to drastically overestimate future IT demands and significantly oversize their IT hardware or DCPI. Taking advantage of the rise in modular DC technology, including close-coupled cooling, to quickly install new capacity only when it is needed both improves the PUE and decreases the final TCO by reducing unnecessary capital expenses. Oversizing often occurs because DC operators do not know, on the whole, how much of their installed IT capacity is being used at any given moment nor how well their DCPI is performing. Installing sensors, meters, and some type of monitoring or analysis software, such as a DCIM solution, is key to knowing when new IT capacity is actually needed as well as identifying opportunities for increased energy efficiencies.

As the majority of DCs are still air cooled, the biggest takeaways for the DoD to improve the efficiency of their DC cooling systems are to focus on improving the airflow management throughout the facility, investigating hot/cold containment technologies, and installing decoupled economizer modes on their cooling systems. Many simple airflow management practices are not followed as the operators do not realize their significant impact on energy efficiency. The next step would be hot/cold air containment strategies, which can be performed with a variety of vendor products, and some exist that are both low cost and high performing. The end result is a significantly more effective cooling system that can either be turned down to consume less energy or accommodate greater IT densities to help meet the increasing IT capability demands across the DoD in a cost-effective manner. Decoupled economizer modes make use of the surrounding environment such that the DC only has to turn on the chiller when absolutely necessary, reducing the operating expenses by improving a DC's annualized PUE. Combined, better airflow management (including containment) and economizer modes represent a proven

and relatively simple way to both increase the efficiency of the DC cooling system and reduce the amount of time it is running. The potential impact is significant as the cost of running a DC cooling system is typically second only to that of the IT equipment itself.

Although the DoD has a significant number of traditional, large-scale DC facilities, they also have a very large number of smaller, tactical systems. Recall that tactical C4 equipment is included in the definition of a DC for the purposes of this report. During this review, multiple technologies stood out for their potential to benefit these unique DoD applications and their feasibility should be investigated. The first is for the DoD to investigate whether the prefabricated modular DCs (both in ISO shipping containers and in smaller configurations) can meet, or be upgraded to meet, military survivability specifications. With significantly faster deployment speeds than traditional IT infrastructure these can potentially be utilized when setting up or expanding FOBs and CPs. Additionally, they could simplify the expansion or modernization of platform DCs, such as on naval vessels or air platforms. Virtualization is another powerful tool for improving the cost effectiveness, and even resiliency, of mobile, tactical DCs. In addition to reducing the amount of physical hardware necessary to field a specific IT or C4 capability, it improves the ability of the DC to compensate for individual system failure and continue to carry out its mission. Virtualization can also have significant benefits for the DoD IT systems that require higher levels of physical security and access control, e.g., DISA Impact Levels 5 & 6 information. It would decrease the difficulty of meeting these requirements by either shrinking down systems to reduce the cost of implementation or increasing the capabilities of systems within existing physical and access control frameworks.

There are a couple of DC cooling technologies that stand out as potentially beneficial for mobile, tactical DC platforms, as well. When C4 capabilities are integrated into a platform whose primary purpose is not C4, there can be a lot of system-level inefficiencies. Installing stand-alone DCs into air and sea platforms means that the vast Free Cooling capabilities of the surrounding high altitude atmosphere or seawater are likely left untapped. Integrating the DC more fully into the platform itself to be able to implement economizer modes for Free Cooling, particularly for long endurance C4ISR platforms, would either drastically reduce the auxiliary power demands the DC draws from the platform or significantly increase the maximum IT payload. Also, although aircraft carriers make use of waste heat from prime power generation systems for desalination, prime power waste heat could be used to power absorption chillers, as well. Either that, or it should be investigated whether the DC waste heat could be used to supplement the prime power generation waste heat and boost the existing systems capabilities. All of these cooling technologies, however, require the platform to be designed in conjunction with the DC on an overall system level; a very significant challenge.

7 References

- [1] Wired Real Estate Group, "Data Center Power Innovation to Reduce TCO," [Online]. Available: <http://wiredre.com/data-center-power-innovation-to-reduce-tco/>. [Accessed 20 August 2015].
- [2] DOE, "Better Buildings Progress Report," DOE, 2015.
- [3] Wired Real Estate Group, "Webinar Series - WiredRE Research Shows Over-Specifying and Over-Provisioning are Top Risks in Cloud, Colo, & Data Center Planning, Investing," 29 September 2014. [Online]. Available: <http://wiredre.com/webinar-series-wiredre-research-shows-specifying-provisioning-top-risks-cloud-colo-data-center-planning-investing/>. [Accessed 20 August 2015].
- [4] DOE, "Better Buildings Initiative Data Center Accelerator," [Online]. Available: <http://www1.eere.energy.gov/buildings/betterbuildings/accelerators/datacenter.html>. [Accessed 30 September 2015].
- [5] Department of Defense, "DoD 2011 Data Center Consolidation Plan & Progress Report," 2011.
- [6] J. Serbu, "DoD now emphasizes cost over quantity in data center consolidation," Federal News Radio, 19 March 2015. [Online]. Available: <http://federalnewsradio.com/defense/2015/03/dod-now-emphasizes-cost-over-quantity-in-data-center-consolidation/>. [Accessed 8 September 2015].
- [7] DISA, "Department of Defense Cloud Computing Security Requirements Guide," 2015.
- [8] Northrop Grumman, *Command Post Platform Rigid Wall Shelter*, 2013.
- [9] U.S. Air Force, "E-8C Joint Stars," 31 March 2005. [Online]. Available: <http://www.af.mil/AboutUs/FactSheets/Display/tabid/224/Article/104507/e-8c-joint-stars.aspx>. [Accessed 8 September 2015].
- [10] Defense Industry Daily, "JSTARS Replacement: Competition Opened Wide," 11 August 2015. [Online]. Available: <https://www.defenseindustrydaily.com/jumped-up-jstars-mprtip-technology-for-ground-surveillance-planes-05156/>. [Accessed 8 September 2015].
- [11] J. Gordon IV, P. A. Wilson, J. Birkler, S. Boraz and G. T. Lee, "Leveraging America's Aircraft Carrier Capabilities: Exploring New Combat and Noncombat Roles and Missions for the U.S. Carrier Fleet," RAND National Defense Research Institute, 2006.
- [12] 42U, "42U Rack Dimensions & Specifications," [Online]. Available: <http://www.42u.com/42U-cabinets.htm>. [Accessed 20 July 2015].
- [13] S. Lowe, "TechRepublic," 5 July 2011. [Online]. Available: <http://www.techrepublic.com/blog/the-enterprise-cloud/the-pros-and-cons-of-tower-rack-and-blade-servers/>. [Accessed 21 7 2015].
- [14] SanDisk, "Data Center Solutions: Systems, software, and hardware solutions you can trust," SanDisk, Milpitas, CA, 2015.
- [15] HP, "HP 10500 Switch Series - Modular Ethernet Switches," [Online]. Available: <http://www8.hp.com/us/en/products/networking-switches/product-detail.html?oid=5117468#!tab=features>. [Accessed 20 August 2015].
- [16] ASHRAE TC 9.9, Thermal Guidelines for Data Processing Environments, 3rd Ed. ed., 2013.

- [17] N. Rasmussen, "Electrical Efficiency Measurement for Data Centers," Schneider Electric, 2011.
- [18] V. Avelar, D. Azevedo and A. French, "PUE: A comprehensive examination of the metric," The Green Grid, 2012.
- [19] ASHRAE TC9.9, PUE: A Comprehensive Examination of the Metric, 2013.
- [20] Lawrence Berkeley National Laboratory Information Technology, "Data Center Utilization Report - August 2011," 15 September 2011. [Online]. Available: <https://commons.lbl.gov/pages/viewpage.action?pageId=76185788>. [Accessed 30 October 2015].
- [21] Y. Sverdlik, "Survey: Industry Average Data Center PUE Stays Nearly Flat Over Four Years," Data Center Knowledge, 02 06 2014. [Online]. Available: <http://www.datacenterknowledge.com/archives/2014/06/02/survey-industry-average-data-center-pue-stays-nearly-flat-four-years/>. [Accessed 21 07 2015].
- [22] Google, "Data Centers Efficiency: How we do it - Measuring Efficiency," Google, [Online]. Available: <http://www.google.com/about/datacenters/efficiency/internal/>. [Accessed 21 07 2015].
- [23] Standard Performance Evaluation Corporation, "SPECpower_ssj2008 Results," [Online]. Available: https://www.spec.org/power_ssj2008/results/. [Accessed 8 June 2015].
- [24] Tier44 Technologies, Inc., "The PAR4 Energy Efficiency Rating System," Tier44, Santa Clara, CA, 2015.
- [25] Upsite Technologies, "CCF Calculator," [Online]. Available: <http://www.upsite.com/resources/ccf-sign-up/>. [Accessed 7 October 2015].
- [26] K. G. Brill and L. Strong, "Cooling Capacity Factor Reveals Stranded Capacity and Data Center Cost Savings," Upsite Technologies, 2013.
- [27] C. Gaskins, "Data Center Efficiency: The Benefits of RCI & RTI (Part 1 of 3)," RF Code, 15 October 2012. [Online]. Available: <http://www.rfcode.com/data-driven-data-center/bid/232990/Data-Center-Efficiency-The-Benefits-of-RCI-RTI-Part-1-of-3>. [Accessed 7 October 2015].
- [28] C. Gaskins, "Data Center Efficiency: The Benefits of RCI & RTI (Part 2 of 3)," RF Code, 22 October 2012. [Online]. Available: <http://www.rfcode.com/data-driven-data-center/bid/233382/Data-Center-Efficiency-The-Benefits-of-RCI-RTI-Part-2-of-3>. [Accessed 7 October 2015].
- [29] J. Neudorfer, "Dehydrating the Data Center: Is your data center thirstier than a golf course?," *Mission Critical*, pp. 6-10, July/August 2015.
- [30] D. Azevedo, C. Belady and J. Pouchet, "Water Usage Effectiveness (WUE): A Green Grid Data Center Sustainability Metric," The Green Grid, 2011.
- [31] N. Rasmussen, "Implementing Energy Efficient Data Centers," Schneider Electric, 2011.
- [32] GSA, DOE FEMP, "Quick Start Guide to Increase Data Center Energy Efficiency," GSA, 2008.
- [33] NREL, "Best Practices Guide for Energy-Efficient Data Center Design," FEMP, 2011.
- [34] D. Bhandarkar, P. Boden, M. Drake, M. Estberg, V. Gill, B. Janous and K. Vaid, "Microsoft's Top 10 Business Practices for Environmentally Sustainable Datacenters," Microsoft Corporation, 2014.
- [35] Energy Star, "12 Ways to Save Energy in Server Rooms and Data Centers," EPA,

- [Online]. Available:
http://www.energystar.gov/index.cfm?c=power_mgt.datacenter_efficiency. [Accessed 28 September 2015].
- [36] Google, "Data Centers Efficiency: How others can do it," [Online]. Available:
<http://www.google.com/about/datacenters/efficiency/external/index.html#best-practices>.
 [Accessed 21 07 2015].
 - [37] ENERGY STAR, "Program Requirements for Computer Servers," EPA, 2013.
 - [38] S. Dale, "Using New Server Telemetry Data to Drive Data Center Excellence," in *Data Center World*, National Harbor, MD, 2015.
 - [39] Packet Power, "Wireless Power Monitoring," [Online]. Available:
<http://www.packetpower.com/products>. [Accessed 23 July 2015].
 - [40] Veris Industries, "Complete Source for Environmental Sensors & Energy Management Solutions," [Online]. Available: <http://www.veris.com/>. [Accessed 13 October 2015].
 - [41] Enlogic, "Enlogic launches energy management solutions in the US," [Online]. Available:
<http://www.enlogic.com/Information/Enlogic-Launches-in-the-U-S/#>. [Accessed 13 October 2015].
 - [42] Methode Electronics, Inc., "PDUs Designed for Higher Density, Increase Reliability, and Continuous Power," [Online]. Available: <http://www.methode.com/data/energy-power-management/pdus.html#.Vh1aZPIVhBd>. [Accessed 13 October 2015].
 - [43] Raritan, "PX Intelligent Rack PDUs," [Online]. Available:
<http://www.raritan.com/products/power-distribution/intelligent-rack-pdus>. [Accessed 13 October 2015].
 - [44] STARLINE, "Critical Power Monitor," [Online]. Available:
<http://www.starlinepower.com/busway/cpm/>. [Accessed 13 October 2015].
 - [45] E+I Engineering Group, "PowerBar BusBar Trunking," [Online]. Available:
<http://www.e-i-eng.com/powerbar-busbar-trunking>. [Accessed 13 October 2015].
 - [46] Emerson Network Power, "Rack Power Distribution," [Online]. Available:
<http://www.emersonnetworkpower.com/en-US/Products/ACPower/PowerDistribution/RackPDU/Pages/Default.aspx>. [Accessed 13 October 2015].
 - [47] Graphical Networks, "netTerrain DCIM: Real-Time Visualization of Data Centers," [Online]. Available: <http://graphicalnetworks.com/netterrain-dcim-2/>. [Accessed 13 October 2015].
 - [48] Methode Electronics, Inc., "DCIM Solutions," [Online]. Available:
http://www.methode.com/data/dcim-solutions.html#.Vh1hl_IVhBd. [Accessed 13 October 2015].
 - [49] CA Technologies, "CA DCIM," [Online]. Available: <http://www.ca.com/us/opscenter/ca-dcim.aspx>. [Accessed 13 October 2015].
 - [50] Raritan, "DCIM Software," [Online]. Available: <http://www.raritan.com/products/dcim-software>. [Accessed 13 October 2015].
 - [51] Emerson Network Power, "Data Center Monitoring: Trellis Platform," [Online]. Available: <http://www.emersonnetworkpower.com/en-US/Brands/Trellis/Pages/default.aspx>. [Accessed 13 October 2015].
 - [52] Sunbird, "The know-it-all DCIM Solution," [Online]. Available:

- <http://www.sunbirddcim.com/>. [Accessed 13 October 2015].
- [53] Manage Engine, "RackBuilder Plus Energy Management," [Online]. Available: <https://www.manageengine.com/data-center-visualization/energy-management.html>. [Accessed 13 October 2015].
- [54] CommScope, "iTRACS Data Center Infrastructure Management," [Online]. Available: [http://www.commscope.com/Solutions/Data-Center-Infrastructure-Management-\(DCIM\)/](http://www.commscope.com/Solutions/Data-Center-Infrastructure-Management-(DCIM)/). [Accessed 13 October 2015].
- [55] E+I Engineering, "Data Centre Infrastructure Management Solution," [Online]. Available: <http://www.e-i-eng.com/energy-management-systems/data-centre-infrastructure-management-solutions>. [Accessed 13 October 2015].
- [56] NetZoomDC, "The most comprehensive DCIM Software Solution," [Online]. Available: <http://www.netzoomdc.com/>. [Accessed 13 Oct 2015].
- [57] Cormant, "Data Center Infrastructure Management," [Online]. Available: <http://cormant.com/solutions/data-center-management-dcim/>. [Accessed 13 October 2015].
- [58] B. Myatt, "DCIM'd if you do and DCIM'd if you don't: failure to implement a DCIM decision may cost you," *Mission Critical Magazine*, pp. 20-22, July/August 2015.
- [59] TileFlow, "Data center CFD modeling software overview," [Online]. Available: <http://tileflow.com/tileflow/data-center-cfd-modeling-software-overview.html>. [Accessed 27 July 2015].
- [60] Electronic Environments, "Data Center Computational Fluid Dynamics (CFD) 3-D Modeling Service," [Online]. Available: <http://www.eecnet.com/Services/Data-Center-Consulting/Computational-Fluid-Dynamics/>. [Accessed 27 July 2015].
- [61] CoolSim, "CoolSim Software Homepage," [Online]. Available: <http://coolsimsoftware.com/>. [Accessed 27 July 2015].
- [62] Future Facilities, "6SigmaDC Overview," [Online]. Available: <http://www.futurefacilities.com/software/6SigmaDCOverview.php>. [Accessed 27 July 2015].
- [63] N. Rasmussen, "Avoiding Costs from Oversizing Data Center and Network Room Infrastructure," Schneider Electric, 2012.
- [64] HP, "Performance Optimized Datacenters," [Online]. Available: <http://www8.hp.com/us/en/products/servers/scalable-systems/pods.html>. [Accessed 23 July 2015].
- [65] SGI, "SGI Announces New ICE Cube Air Modular Data Center," 6 December 2010. [Online]. Available: http://www.sgi.com/company_info/newsroom/press_releases/2010/december/ice_cube_air.html. [Accessed 23 July 2015].
- [66] Baselayer Technology, "Baselayer Products," [Online]. Available: <https://www.baselayer.com/products/>. [Accessed 3 August 2015].
- [67] Dell, "Dell Modular Data Center," [Online]. Available: <http://www.dell.com/learn/us/en/04/solutions/modular-data-center>. [Accessed 23 July 2015].
- [68] Silent-Aire, "Modular IT Solutions," [Online]. Available: <http://www.silent-aire.com/products/modularitsolutions>. [Accessed 23 July 2015].

- [69] Microsoft Corporation, "Real Impact for a Better Environment: UNEP and Microsoft Showcasing Design and Green Technology," Microsoft Corporation, 2011.
- [70] Supernap, "The Rob Roy Factor," Switch, [Online]. Available: <http://www.supernap.com/rob-roy-factor.html>. [Accessed 23 July 2015].
- [71] Schneider Electric, "Prefabricated Data Center Modules - Fast, Flexible, Predictable," [Online]. Available: <http://www.schneider-electric.com/solutions/ww/en/app/4664756-data-centre-systems/633790233-prefabricated-data-center-modules-fast-flexible-predictable>. [Accessed 23 July 2015].
- [72] Schneider Electric, "Micro Data Centers," [Online]. Available: <http://www.schneider-electric.com/solutions/ww/en/app/4664756-data-centre-systems/843297067-micro-data-centers>. [Accessed 8 October 2015].
- [73] Panduit, "Pre-Configured Micro Data Center," [Online]. Available: <http://www.panduit.com/en/solutions/industrial-solutions/offerings/micro-data-center>. [Accessed 8 October 2015].
- [74] Rittal, "Micro Data Center," [Online]. Available: http://www.rittal.com/de_de/micro_data_center/?l=en. [Accessed 8 October 2015].
- [75] Black Box, "Climate Controlled Cabinets," [Online]. Available: <http://www.blackbox.com/Store/Results.aspx/Cabinets-Racks/Full-Size-IT-Cabinets/Climate-Controlled-Cabinets/n-4294954894>. [Accessed 8 October 2015].
- [76] Wired Real Estate Group, "Modular Data Center Projects on the Rise," [Online]. Available: <http://wiredre.com/modular-data-center-projects-on-the-rise/>. [Accessed 20 August 2015].
- [77] D. Bouley and W. Torell, "Prefabricated Power and Cooling Modules for Data Centers," Schneider Electric, 2014.
- [78] W. Torell, "TCO Analysis of a Traditional Data Center vs. a Scalable, Prefabricated Data Center," Schneider Electric, 2014.
- [79] BladeRoom, "BladeRoom Modular Data Centres," [Online]. Available: <http://www.bladeroom.com/>. [Accessed 3 August 2015].
- [80] N. Rasmussen, "Electrical Efficiency Modeling for Data Centers," Schneider Electric, 2011.
- [81] AER Worldwide, "AER Worldwide: The global leader in E-Recycling, IT Asset Management, Data Security & Value Recover," [Online]. Available: <http://www.aerworldwide.com/>. [Accessed 29 September 2015].
- [82] Potomac eCycle, "Potomac eCycle: Home," [Online]. Available: <http://potomacecycle.com/>. [Accessed 29 September 2015].
- [83] L. A. Barroso and U. Holzle, "The case for energy-proportional computing," *IEEE Computer*, vol. 40, pp. 33-37, 2007.
- [84] Energy Star, "Enterprise Servers ENERGY STAR Certification Program," [Online]. Available: <http://www.energystar.gov/products/certified-products/detail/7620/partners>. [Accessed 30 September 2015].
- [85] D. Garday, "Reducing Data Center Energy Consumption with Wet Side Economizers," Intel, 2007.
- [86] Pure Storage, "The Remarkable ROI of All-Flash Arrays," MIT Technology Reivew Custom , 2015.

- [87] Seagate, "Desktop SSHD," [Online]. Available: <http://www.seagate.com/internal-hard-drives/solid-state-hybrid/desktop-solid-state-hybrid-drive/>. [Accessed 30 October 2015].
- [88] VMware, "Server Consolidation - Overview," [Online]. Available: <http://www.vmware.com/consolidation/overview.html>. [Accessed 4 August 2015].
- [89] Microsoft, "Virtualization," [Online]. Available: <http://www.microsoft.com/en-us/server-cloud/solutions/virtualization.aspx>. [Accessed 28 July 2015].
- [90] VMWare, "The Software-Defined Data Center," [Online]. Available: <http://www.vmware.com/software-defined-datacenter/index.html>. [Accessed 28 July 2015].
- [91] E. Hsu and J. Townsend, "VMware Software-Defined Data Center - Capabilities and Outcomess," VMware, 2015.
- [92] S. Sheffield and J. Mandelbaum, "The Benefits of VMware's vCenter Operations Management Suite: Quantifying the Incremental Value of the vCenter Operations Management Suites for vSphere Customers," Management Insight Technologies, 2012.
- [93] VMware, "VMware distributed power management concepts and use," VMware, 2010.
- [94] D. M. Shinohara and E. Symonds, "Data Center Consolidation Energy Efficiency Pilot," Air Force Space Command, 2012.
- [95] L. Martin, "Energy Savings Performance Contracts," U.S. Air Force Headquarters, 2012.
- [96] 42U, "Hot Aisle Containment Strategies," [Online]. Available: <https://blog.mozilla.org/it/2011/11/22/project-scl3-containment-hot-or-cold/>. [Accessed 10 August 2015].
- [97] C. Jaggars and T. North, "Airflow Containment: The Road to Data Center Efficiency," *BICSI News Magazine*, May/June 2010.
- [98] D. Garday and D. Costello, "Air-Cooled High-Performance Data Centers: Case Studies and Best Methods," Intel, 2006.
- [99] S. Greenberg, "Variable-Speed Fan Retrofits for Computer-Room Air Conditioners," DOE Federal Energy Management Program, 2013.
- [100] P. Matte, "Airflow and Pressure Behavior: A Practical vs Theoretical Examination," in *Data Center World*, National Harbor, MD, 2015.
- [101] D. Kennedy, "The 'Small' Data Center Efficiency Potential," in *Data Center World*, National Harbor, MD, 2015.
- [102] Tate Airflow, "Air Sealing Grommets," [Online]. Available: <http://tateinc.com/data-center/accessories/air-sealing-grommets>. [Accessed 7 October 2015].
- [103] SubZero Engineering, "The Cube," [Online]. Available: <http://www.subzeroeng.com/products/airflow-management/cube>. [Accessed 7 October 2015].
- [104] Google, "Data Centers Efficiency: How we do it - Temperature control," [Online]. Available: <http://www.google.com/about/datacenters/efficiency/internal/#temperature>. [Accessed 27 July 2015].
- [105] Simplex, "Simplex AirBlock Expandable Filler Panels," [Online]. Available: <http://www.simplexstripdoors.com/dcexpandable.htm>. [Accessed 8 October 2015].
- [106] SubZero Engineering, "Velocity Adjuster," [Online]. Available: <http://www.subzeroeng.com/products/airflow-management/velocity-adjustor>. [Accessed 7 October 2015].

- [107] L. Strong, "Reducing Room-Level Bypass Airflow Creates Opportunities to Improve Cooling Capacity and Operating Costs," Upsite Technologies, 2013.
- [108] 42U, "Data Center Containment: Chimney Systems," [Online]. Available: <http://www.42u.com/data-center-containment/chimney-systems>. [Accessed 27 July 2015].
- [109] I. Seaton, "The Biggest Opportunity for Data Center Energy Savings," Chatsworth Products, Inc., 2009.
- [110] Google, "Green Blog - Cooling the cloud: A look inside Google's Hot Huts," 23 October 2012. [Online]. Available: <http://googlegreenblog.blogspot.com/2012/10/cooling-cloud-look-inside-googles-hot.html>. [Accessed 27 July 2015].
- [111] 42U, "Chimney Systems," [Online]. Available: <http://www.42u.com/data-center-containment/chimney-systems>. [Accessed 10 August 2015].
- [112] Cirrascale, "Patented Vertical Cooling Technology," [Online]. Available: http://www.cirrascale.com/company_verticalcoolingtechnology.aspx. [Accessed 27 July 2015].
- [113] Simplex, "Data Center Containment Curtains," [Online]. Available: <http://www.simplexstripdoors.com/dcairblock.html>. [Accessed 8 October 2015].
- [114] SubZero Engineering, "Data Center Curtain Containment," [Online]. Available: <http://www.subzeroeng.com/products/data-center-curtain-containment>. [Accessed 8 October 2015].
- [115] 42U, "Data Center Containment Curtains," [Online]. Available: <http://www.42u.com/data-center-containment/curtain-systems>. [Accessed 27 July 2015].
- [116] Upsite Technologies, "Introducing our award-winning containment solution," [Online]. Available: <http://www.upsite.com/modularcontainment/>. [Accessed 27 July 2015].
- [117] Simplex, "Data Center Doors and End-of-Aisle Solutions," [Online]. Available: <http://www.simplexstripdoors.com/dccdoors.html>. [Accessed 8 October 2015].
- [118] B. Long and L. Strong, "AisleLok Modular Containment vs. Legacy Containment: A comparative CFD study of IT inlet temperatures and fan energy savings," Upsite Technologies, 2015.
- [119] Upsite Technologies, "AisleLok Modular Containment: The industry's first modular containment solution designed for quick and easy installation," 2014.
- [120] L. Strong, "Myths of Data Center Containment: What's True and What's Not," in *Data Center World*, National Harbor, MD, 2015.
- [121] Google, "Google's Hamina Data Center," Google Green, 23 May 2011. [Online]. Available: <https://www.youtube.com/watch?v=VChOEvKicQQ>. [Accessed 27 July 2015].
- [122] Google for Work, "Inside a Google Data Center," 16 December 2014. [Online]. Available: <https://www.youtube.com/watch?v=XZmGGAbHqa0>. [Accessed 10 August 2015].
- [123] TelecityGroup, "TelecityGroup opens new state-of-the-art data centre in Paris," 22 January 2010. [Online]. Available: <http://www.telecitygroup.com/our-company/news/2010/telecitygroup-opens-new-state-of-the-art-data-centre-in-paris.htm>. [Accessed 22 July 2015].
- [124] R. Mahdavi, "Case Study: Innovative Energy Efficiency Approaches in NOAA's Environmental Security Computing Center in Fairmont, West Virginia," DOE EERE FEMP, 2014.

- [125] The Green Grid, "Library and Tools," [Online]. Available: <http://www.thegreengrid.org/library-and-tools.aspx>. [Accessed 30 Oct 2015].
- [126] J. Niemann, J. Bean and V. Avelar, "Economizer Modes of Data Center Cooling Systems," Schneider Electric, 2011.
- [127] STULZ, "STULZ Air-Side Economizers," [Online]. Available: <http://www.stulz-ats.com/solutions/stulz-air-side-economizers/>. [Accessed 30 September 2015].
- [128] Munters, "Oasis Indirect Evaporative Cooling Systems - Americas," [Online]. Available: <https://www.munters.com/en/munters/products/coolers--humidifiers/oasis-indirect-evaporative-coolers-iec/>. [Accessed 39 September 2015].
- [129] KyotoCooling, "KyotoCooling Theorem of Operation," 2013. [Online]. Available: <http://www.kyotocooling.com/en/theorem%20of%20operation.html>. [Accessed 27 July 2015].
- [130] K. Dunnivant, "Indirect Air-Side Economizer Cycle Data Center Heat Rejection," *ASHRAE Journal*, pp. 44-54, March 2011.
- [131] T. J. J. Lawrence, "Data Center Fresh Air Cooling with Heat Wheels," in *Data Center World*, Orlando, 2014.
- [132] KyotoCooling, "Why KyotoCooling?," [Online]. Available: <http://www.kyotocooling.com/en/index.html>. [Accessed 11 August 2015].
- [133] Mee Industries, Inc., "Using Outside Air and Evaporative Cooling to Increase Data Center Efficiency," 2014.
- [134] Google, "Green Blog - Getting the heat out: A look at evaporative cooling," 19 November 2012. [Online]. Available: <http://googlegreenblog.blogspot.com/2012/11/getting-heat-out-look-at-evaporative.html>. [Accessed 27 July 2015].
- [135] M. Ton, B. Fortenbery and W. Tschudi, "DC Power for Improved Data Center Efficiency," Lawrence Berkeley National Laboratory, 2008.
- [136] Intel, "Intel Core i7-900 Desktop Processor Extreme Edition Series and Intel Core i7-900 Desktop Processor Series," 2010.
- [137] M. Murrill and B. Sonnenberg, "Evaluating the Opportunity for DC Power in the Data Center," Emerson Network Power Energy Systems, North America, Inc., 2010.
- [138] S. Harris, "Microsoft Reinvents Datacenter Power Backup with New Open Compute Project Specification," Microsoft Server and Cloud Platform Team, 10 March 2015. [Online]. Available: <http://blogs.technet.com/b/server-cloud/archive/2015/03/10/microsoft-reinvents-datacenter-power-backup-with-new-open-compute-project-specification.aspx>. [Accessed 12 August 2015].
- [139] Google, "Data Centers Efficiency: How we do it - Servers," [Online]. Available: <http://www.google.com/about/datacenters/efficiency/internal/#servers>. [Accessed 22 07 2015].
- [140] Emerson Network Power, "NetSure ITM with eSure Technology," [Online]. Available: <http://www.emersonnetworkpower.com/en-US/Products/DCPower/DC%20UPS/Pages/NetSureITM48VDCUPS70-280kW.aspx>. [Accessed 12 August 2015].
- [141] The Green Grid, "Quantitative Efficiency Analysis of Power Distribution Configurations for Data Centers," 2008.
- [142] T. Evans, "Humidification Strategies for Data Centers and Network Rooms," Schneider

- Electric, 2011.
- [143] Mee Industries, "Data center uses MeeFog system to keep operating through a heat wave," 2012.
 - [144] 42U, "Close-Coupled Cooling," [Online]. Available: <http://www.42u.com/cooling/close-coupled-water-cooled.htm>. [Accessed 23 July 2015].
 - [145] J. Capes, "The Economics of Close-Coupled Cooling in the Data Center," in *INTEROP*, Las Vegas, 2011.
 - [146] J. Cermak and M. Ivanovich, "Fan Efficiency Requirements for Standard 90.1-2013," *ASHRAE Journal*, April 2013.
 - [147] OptiCool Technologies, "OptiCool Overview," [Online]. Available: <http://www.opticooltechnologies.com/opticool-overview/>. [Accessed 21 August 2015].
 - [148] IBM Zurich, "Zero-emission datacenter: direct use of waste heat to minimize carbon-dioxide emission," [Online]. Available: <http://www.zurich.ibm.com/st/energy/zeroemission.html>. [Accessed 27 July 2015].
 - [149] PNNL, "Hot Topic - the Problem of Cooling Supercomputers," [Online]. Available: http://www.pnl.gov/computing/resources/esdc/1_Cooling.stm. [Accessed 27 July 2015].
 - [150] T. Brunschweiler, B. Smith, E. Ruetsche and B. Michel, "Toward zero-emission data center through direct reuse of thermal energy," *IBM Journal of Research and Development*, vol. 53, no. 3, pp. 11:1-11:13, 2009.
 - [151] K. Ebrahimi, G. F. Jones and A. S. Fleischer, "A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery options," *Renewable and Sustainable Energy Reviews*, vol. 31, pp. 622-638, 2014.
 - [152] R. Miller, "Data Center Used to Heat Swimming Pool," Data Center Knowledge, 2 April 2008. [Online]. Available: <http://www.datacenterknowledge.com/archives/2008/04/02/data-center-used-to-heat-swimming-pool/>. [Accessed 22 July 2015].
 - [153] R. Miller, "Telehouse to Heat Homes at Docklands," Data Center Knowledge, 15 April 2009. [Online]. Available: <http://www.datacenterknowledge.com/archives/2009/04/15/telehouse-to-heat-homes-at-docklands/>. [Accessed 22 July 2015].
 - [154] T. Virki, "Cloud computing goes green underground in Finland," Reuters, 30 November 2009. [Online]. Available: <http://www.reuters.com/article/2009/11/30/idUSGEE5AS01D20091130>. [Accessed 22 July 2015].
 - [155] DOE Office of Energy Efficiency & Renewable Energy, "Absorption Cooling Basics," 16 August 2013. [Online]. Available: <http://energy.gov/eere/energybasics/articles/absorption-cooling-basics>. [Accessed 19 August 2015].
 - [156] Energy Solutions Center, "Gas-Fired Air Conditioning Equipment - How Absorption Cooling Works," [Online]. Available: http://www.gasairconditioning.org/absorption_how_it_works.htm. [Accessed 19 August 2015].
 - [157] EPA Combined Heat and Power Partnership, "Combined Heat and Power: Energy Savings and Energy Reliability for Data Centers," EPA, 2008.
 - [158] J. Garthwaite, "KPMG Captures Heat for Data Center Cooling," National Geographic

- News, 30 October 2011. [Online]. Available: <http://news.nationalgeographic.com/news/energy/2011/10/111028-kpmg-chp-data-center-cooling/>. [Accessed 22 July 2015].
- [159] Data Center Knowledge, "Microturbines provide heat, power, and cooling," [Online]. Available: <http://www.datacenterknowledge.com/inside-ibms-greenest-data-center/microturbines-provide-heat-power-and-cooling/>. [Accessed 22 July 2015].
 - [160] Syracuse University, "Spotlight: IBM Green Data Center," [Online]. Available: <http://www.syr.edu/partnerships/corporations/green-data-center.html>. [Accessed 19 August 2015].
 - [161] Syracuse University, "As the saying goes, information is power. The Green Data Center at SU," [Online]. Available: https://www.syr.edu/greendatacenter/GDC_facts.pdf. [Accessed 19 August 2015].
 - [162] Y. Sverdlik, "eBay's gas-powered data center to reclaim waste-heat energy," Datacenter Dynamics, 26 September 2013. [Online]. Available: <http://www.datacenterdynamics.com/news/ebays-gas-powered-data-center-to-reclaim-waste-heat-energy/82414.fullarticle>. [Accessed 22 July 2015].
 - [163] R. Miller, "eBay Goes Live With its Bloom-Powered Data Center," Data Center Knowledge, 26 September 2013. [Online]. Available: <http://www.datacenterknowledge.com/archives/2013/09/26/ebay-goes-live-with-its-bloom-powered-data-center/>. [Accessed 17 February 2015].
 - [164] H. Li, N. Russell, V. Sharifi and J. Swithenbank, "Techno-economic feasibility of absorption heat pumps using wastewater as the heating source for desalination," *Desalination*, vol. 281, pp. 118-127, 2011.
 - [165] Trevi Systems, "Technology: Benefits of the Trevi System," [Online]. Available: <http://trevisystems.com/technology/benefits-of-the-trevi-system/>. [Accessed 17 August 2015].
 - [166] Trevi Systems, "Technology: System Process," [Online]. Available: <http://trevisystems.com/technology/system-process/>. [Accessed 17 August 2015].
 - [167] 3M, "3M Fluorinert Electronic Fluids," [Online]. Available: http://solutions.3m.com/wps/portal/3M/en_US/Electronics_NA/Electronics/Products/Electronics_Product_Catalog/~Electronic-Liquids?N=8704931&rt=r3. [Accessed 21 August 2015].
 - [168] Green Revolution Cooling, "The CarnotJet System," [Online]. Available: <http://www.grcooling.com/carnotjet/>. [Accessed 15 October 2015].
 - [169] LiquidCool Solutions, "Rack Servers," [Online]. Available: <http://www.liquidcoolsolutions.com/rack-servers/>. [Accessed 15 October 2015].
 - [170] Iceotope, "Products," [Online]. Available: <http://www.iceotope.com/#products>. [Accessed 15 October 2015].
 - [171] Allied Control, "Immersion Cooling," [Online]. Available: www.allied-control.com. [Accessed 15 October 2015].
 - [172] Manage Engine, "Energy Management," [Online]. Available: <https://www.manageengine.com/data-center-visualization/energy-management.html>. [Accessed 14 September 2015].
 - [173] Lawrence Livermore National Laboratory, "Energy Efficient High Performance Computing Working Group," [Online]. Available: <https://eehpcwg.llnl.gov/>. [Accessed 14

September 2015].

- [174] T. R. Scogland, C. P. Steffen, T. Wilde, F. Parent, S. Coghlan, N. Bates, W.-c. Feng and E. Strohmaier, "A Power-Measurement Methodology for Large-Scale, High-Performance Computing," in *5th ACM/SPEC International Conference on Performance Engineering (ICPE)*, Dublin, Ireland, 2014.
- [175] Tier44, "EM/8 Runbook Automation," [Online]. Available: <http://www.tier44.com/t44/?CID=gen1300>. [Accessed 8 October 2015].
- [176] TSO Logic, "Out Solution," [Online]. Available: <http://tsologic.com/data-center-power-management/>. [Accessed 15 October 2015].
- [177] C. Pfeiffer, "A Case Study: Ways a Scheduling System can Maximize Energy Efficiency," in *Data Center World*, National Harbor, MD, 2015.
- [178] Intel, "CoFluent - Data Center Performance Modeling," [Online]. Available: <http://www.intel.com/content/www/us/en/cofluent/cofluent-solutions-data-center-modeling.html>. [Accessed 14 September 2015].
- [179] P. Desnoyers, T. Wood, P. Shenoy, R. Singh, S. Patil and H. Vin, "Modellus: Automated modeling of complex internet data center applications," *ACM Transactions on the Web (TWEB)*, vol. 6, no. 2, pp. 8:1 - 8:29, 2012.
- [180] W. Goetzler, R. Zogg, J. Young and C. Johnson, "Energy Savings Potential and RD&D Opportunities for Non-Vapor-Compression HVAC Technologies," US DOE Office of Energy Efficiency and Renewable Energy Building Technologies, 2014.
- [181] ARPA-e, "Climate Control Systems of the Future," *International Innovation*, 8 Feb 2013.
- [182] GE, "Not your average fridge magnet - these high-tech magnets will keep your butter (and beer) cold," 7 February 2014. [Online]. Available: <http://www.gereports.com/post/75911607449/not-your-average-fridge-magnet>. [Accessed 14 September 2015].
- [183] Ames Laboratory, *Room-Temperature Magnetic Refrigeration: An Ames Laboratory Innovation Could Transform Cooling Devices*.
- [184] X.-C. Zhong, Z.-W. Liu, D.-C. Zeng, K. A. J. Gschneider and V. K. Pecharsky, "Magnetocaloric effect of Pr₂Fe₁₇-xMnx alloys," *Rare Metals*, vol. 33, no. 5, pp. 552-555, 2014.
- [185] X. Li, H. Gu, X. Qian and Q. Zhang, "Compact Cooling Devices Based on Giant Electrocaloric Effect Dielectrics," in *IEEE ITherm Conference*, 2012.
- [186] S. Harris, "Electrocaloric material could lead to more efficient refrigerators," Centaur Communications Ltd, 13 December 2013. [Online]. Available: <http://www.theengineer.co.uk/channels/design-engineering/news/electrocaloric-material-could-lead-to-more-efficient-refrigerators/1017688.article>. [Accessed 14 September 2015].
- [187] I. Sauciuc, "Piezo actuators for electronics cooling," *Electronics Cooling*, 1 February 2007.
- [188] S. Wait, "Ultrasonic piezoelectric fans for microelectronics cooling," Purdue University Thesis-MS, 2004.
- [189] N. Zuckerman and N. Lior, "Jet Impingement Heat Transfer: Physics, Correlations, and Numerical Modeling," *Advances in Heat Transfer*, vol. 39, pp. 565-631, 2006.
- [190] Advanced Thermal Solutions, Inc., "What is Jet Impingement Cooling and How is it

- applied for Thermal Management of Electronics (Part 1 of 2)," 13 September 2010. [Online]. Available: <https://heatsinks.wordpress.com/2010/09/13/what-is-jet-impingement-cooling-and-how-is-it-applied-for-thermal-management-of-electronics-part-1-of-2/>. [Accessed 15 September 2015].
- [191] Advanced Thermal Solutions, Inc., "What is Jet Impingement Cooling and How is it applied for Thermal Management of Electronics (Part 2 of 2)," 14 September 2010. [Online]. Available: <https://heatsinks.wordpress.com/2010/09/14/what-is-jet-impingement-cooling-and-how-is-it-applied-for-thermal-management-of-electronics-part-2-of-2/>. [Accessed 15 September 2015].
- [192] Advanced Thermal Solutions, Inc., "Jet Impingement for thermal management: a practical approach to using it in 1U Server Chassis," 13 July 2010. [Online]. Available: <https://heatsinks.wordpress.com/2010/07/13/jet-impingement-for-thermal-management-a-practical-approach-to-using-it-in-1u-server-chassis/>. [Accessed 15 September 2015].
- [193] Advanced Thermal Solutions, Inc., "Single Phase Micro Jet: Arrayed Impingement," *Qpedia*, vol. V, no. II, pp. 10-12, February 2011.
- [194] S. Werner, "FlexTiles: self adaptive heterogeneous manycore based on Flexible Tiles," Thales Research & Technology, [Online]. Available: <http://flextiles.eu/WordPress3/>. [Accessed 15 September 2015].
- [195] Nano-Tera, "CMOSAIC," [Online]. Available: <http://www.nano-tera.ch/projects/67.php#desc>. [Accessed 15 September 2015].
- [196] AMD, "What is Heterogeneous System Architecture (HSA)?," [Online]. Available: <http://developer.amd.com/resources/heterogeneous-computing/what-is-heterogeneous-system-architecture-hsa/>. [Accessed 15 September 2015].
- [197] G. Kyriazis, "Heterogeneous System Architecture: A Technical Review," AMD, 2012.
- [198] J. Watson and G. Castro, "High-temperature electronics pose design and reliability challenges," *Analog Dialog*, April 2012.
- [199] Honeywell Aerospace, "High Temperature HTMOS Products," [Online]. Available: <https://aerospace.honeywell.com/products/sensors/non-inertial-sensors/high-temperature-htmos-products>. [Accessed 16 September 2015].
- [200] CISSOID, "High Temperature and Extended Lifetime Semiconductors," [Online]. Available: <http://www.cissoid.com/>. [Accessed 16 September 2015].
- [201] G. Beheim, P. G. Neudeck and D. J. Spry, "High temperature SiC electronics: update and outlook," in *NASA Glenn Research Center Propulsion Controls and Diagnostics Workshop*, Cleveland, OH, 2012.
- [202] MIT Technology Review, "The Fantastical Promise of Reversible Computing," 25 January 2011. [Online]. Available: <http://www.technologyreview.com/view/422511/the-fantastical-promise-of-reversible-computing/>. [Accessed 16 September 2015].
- [203] N. Takeuchi, Y. Yamanashi and N. Yoshikawa, "Reversible logic gate using adiabatic superconducting devices," *Scientific Reports*, vol. 4, pp. 1-4, 2014.